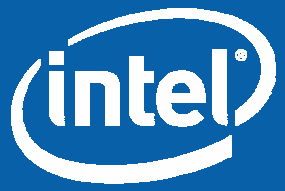

Why Multicore?

- and implications to users and Industry -

Dr. David Barkai, Intel

TECC 2006
October 19, 2006

Outline

- The HPC market
  - And how Intel views it, why it is important
- Why Multicore
- What Intel is doing for HPC
- Implications to you – the users

The HPC Market: The New Realities

- Major market growth over the last 4 years
  - 94% growth since 2002
  - 24% growth in 2005 – Now $9.2 billion a year
- Clusters have been a disruptive force
  - 1/3 of the market in 2004
  - Now close to 1/2 of the market
  - Caused a growth revolution, not a decline
- Capability market transition continues
  - Down -13% since 2002
- Strong growth at the lower end of the market
  - Workgroup up 200%, Departmental up 155%, Divisional up 84% since 2002

Source: IDC 2006

1H06 HPTC Overview

- Market continues to see strong growth
  - 11% Revenue growth 1H06 vs. 1H05 for $4.7B
  - 24% Unit growth 1H06 vs. 1H05 for 113,000 shipments
- All Capacity segments showed growth
  - Departmental -- 32% revenue growth
  - Workgroup -- 16% revenue growth
  - Enterprise and Divisional -- 4% and 3% respectively
- Capability market transition continues
  - Down -43% in revenue
  - Traditionally sees significant variability between quarters
- Clusters continue to gain share
  - 45% Revenue growth 1H06 vs. 1H05
  - 54% Share of revenue

Source: IDC 2006

HPC Segment Revenue Share

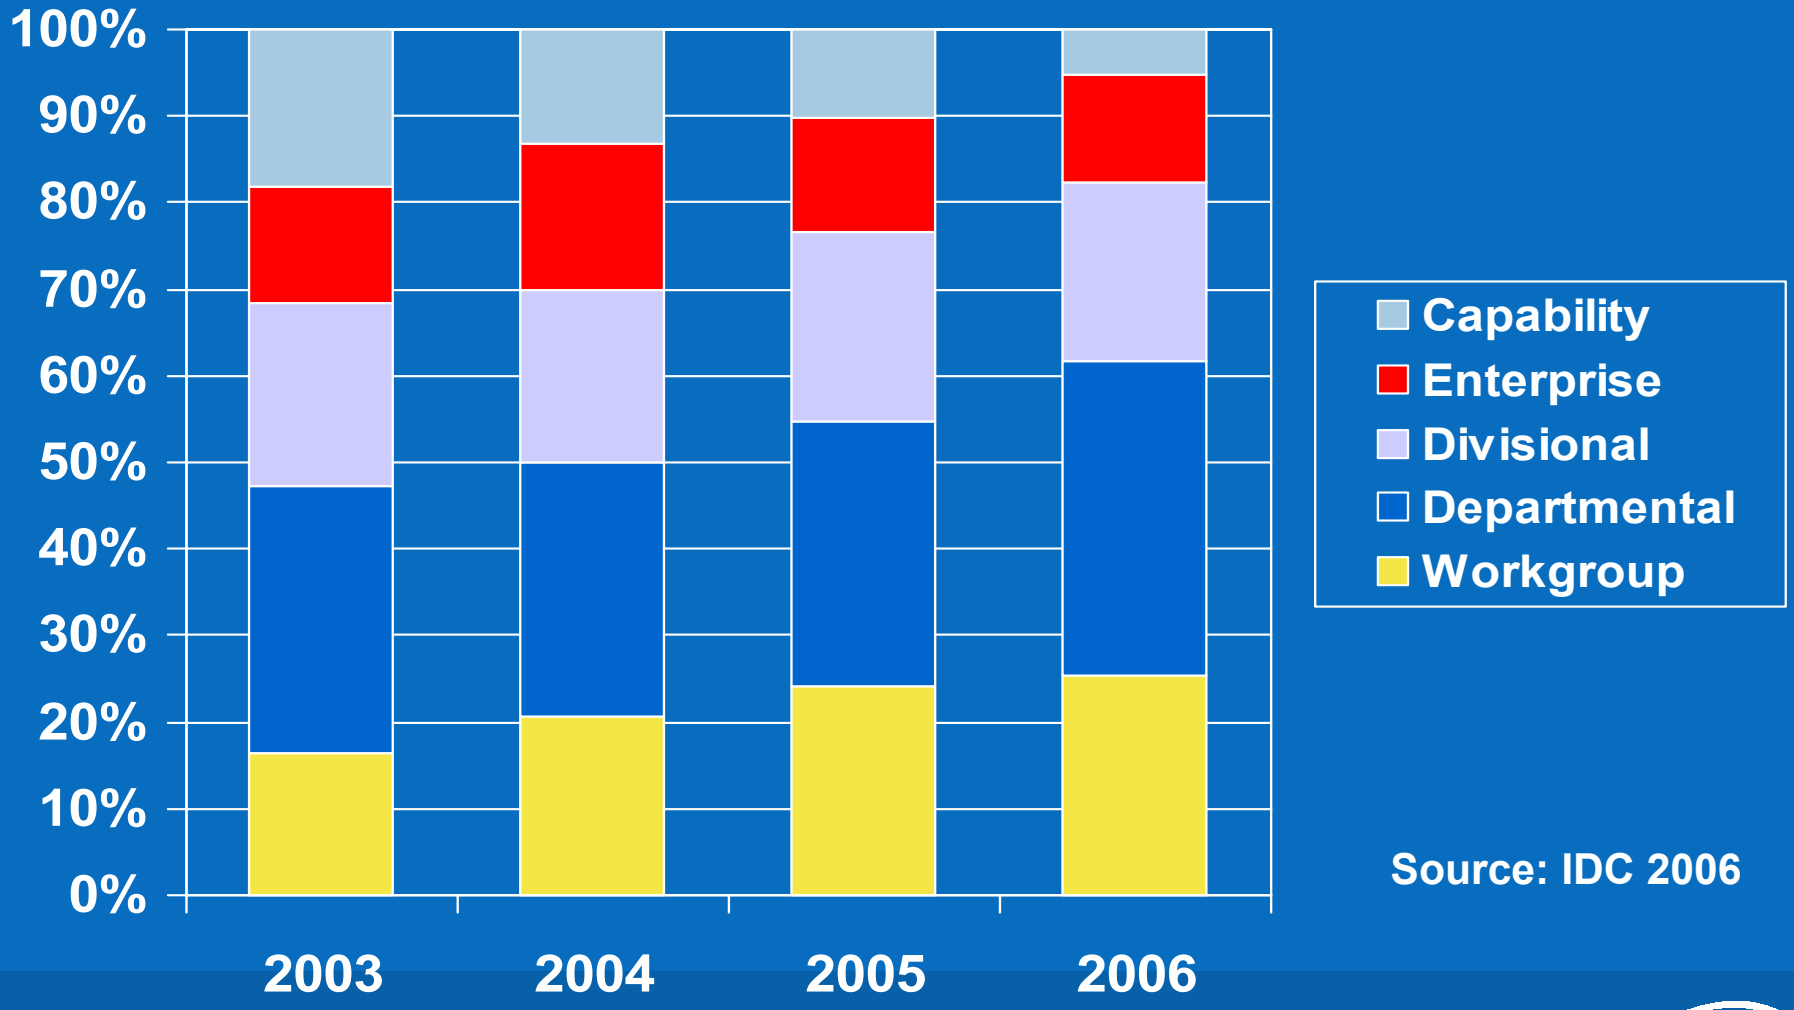

Source: IDC 2006

Application/Industry Segments – 2005 Preliminary Estimates

	2005
Bio-Sciences	$1,198,416
CAE	$1,296,509
Chemical Engineering	$310,197
DCC & Distribution	$144,105
Economics/Financial	$317,344
EDA	$631,773
Geosciences and Geo-engineering	$306,423
Mechanical Design and Drafting	$181,000
Defense	$1,371,911
Government Lab	$1,141,305
Software Engineering	$25,901
Technical Management	$305,883
University/Academic	$1,653,557
Weather	$291,904
Other	$21,792
Total Revenue	$9,198,020

Source: IDC, 2006

HPC Revenue Share by Processor Type

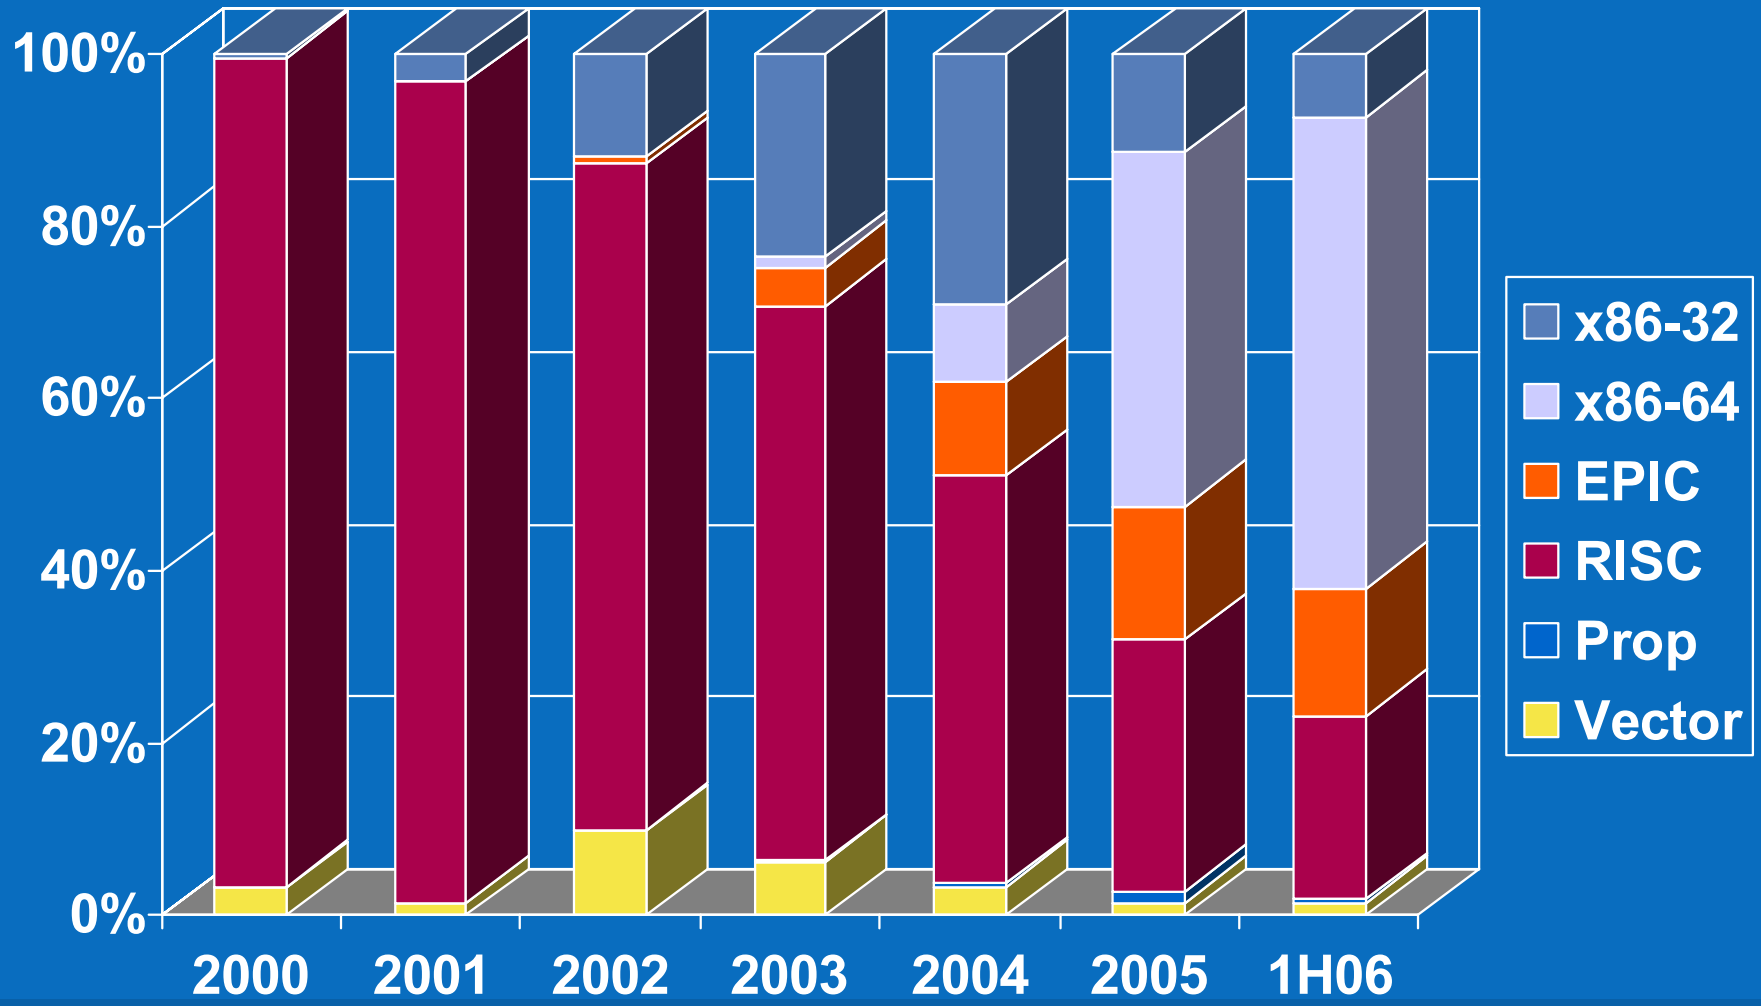

Prop = Proprietary (e.g., FPGA, BlueGene, Cell)

Source: IDC 2006

Growth In Clusters

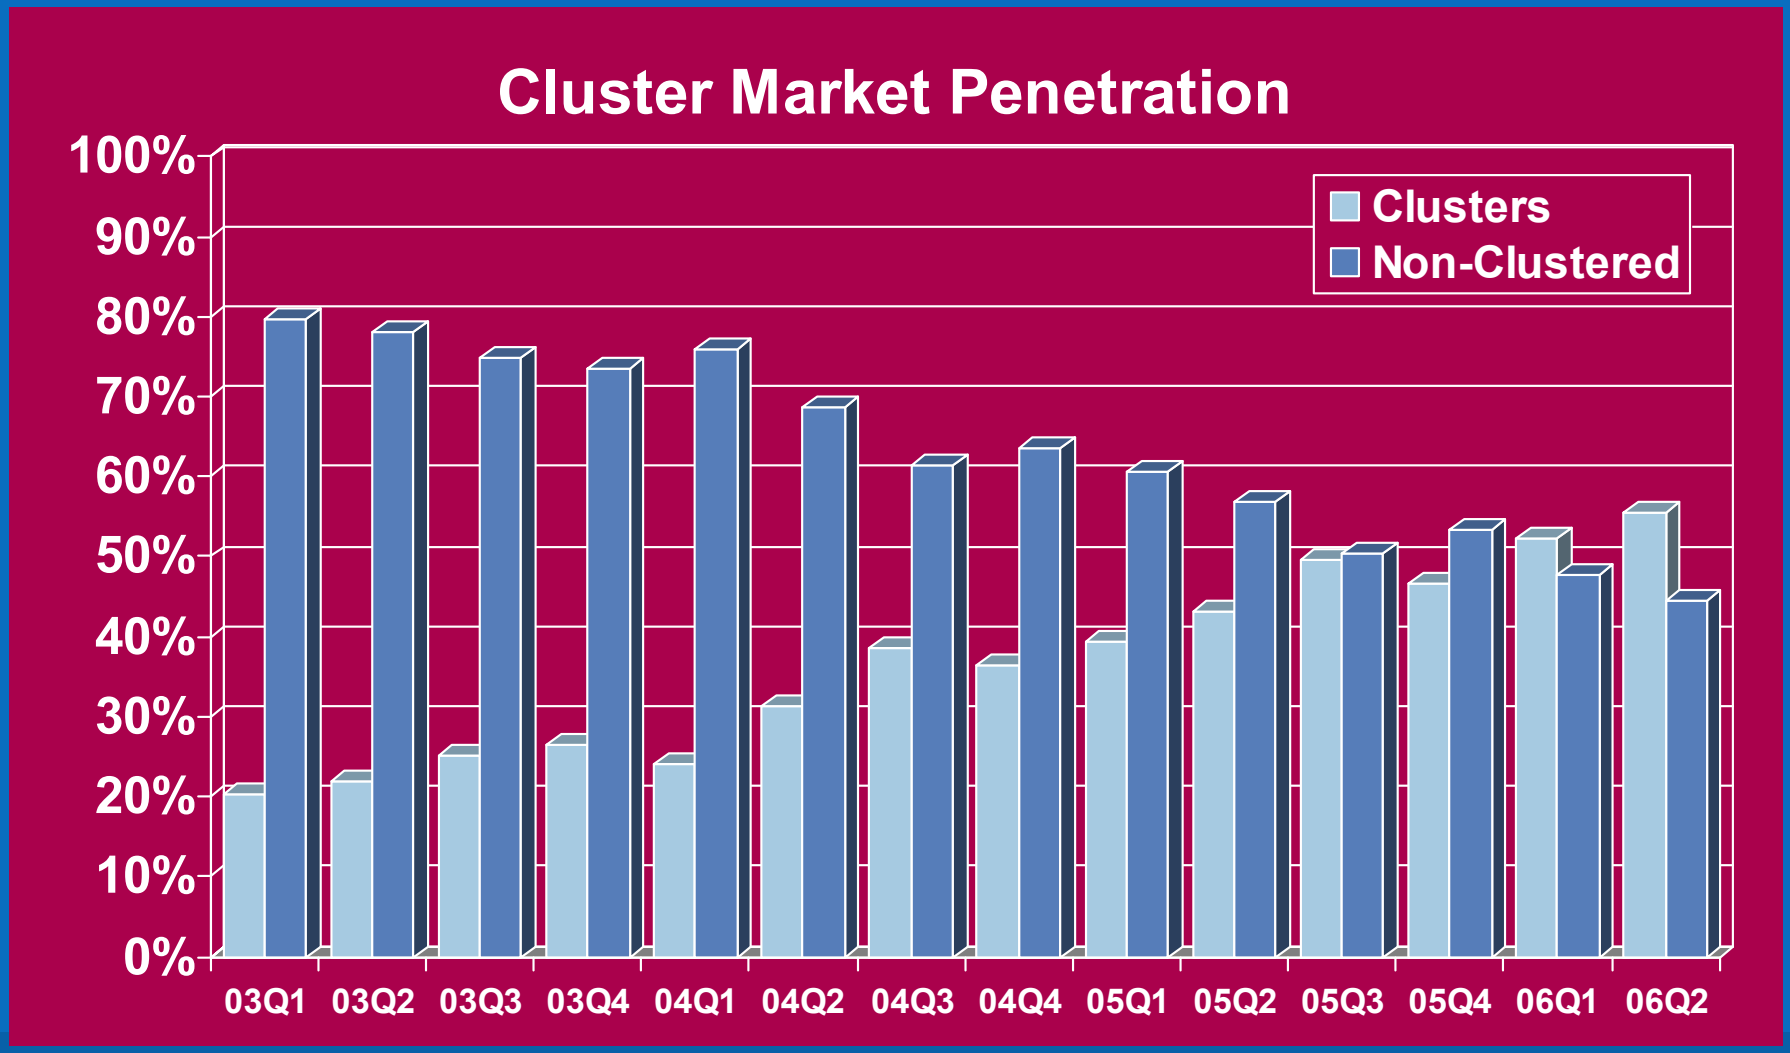

HPC Clusters Revenue Share by Processor

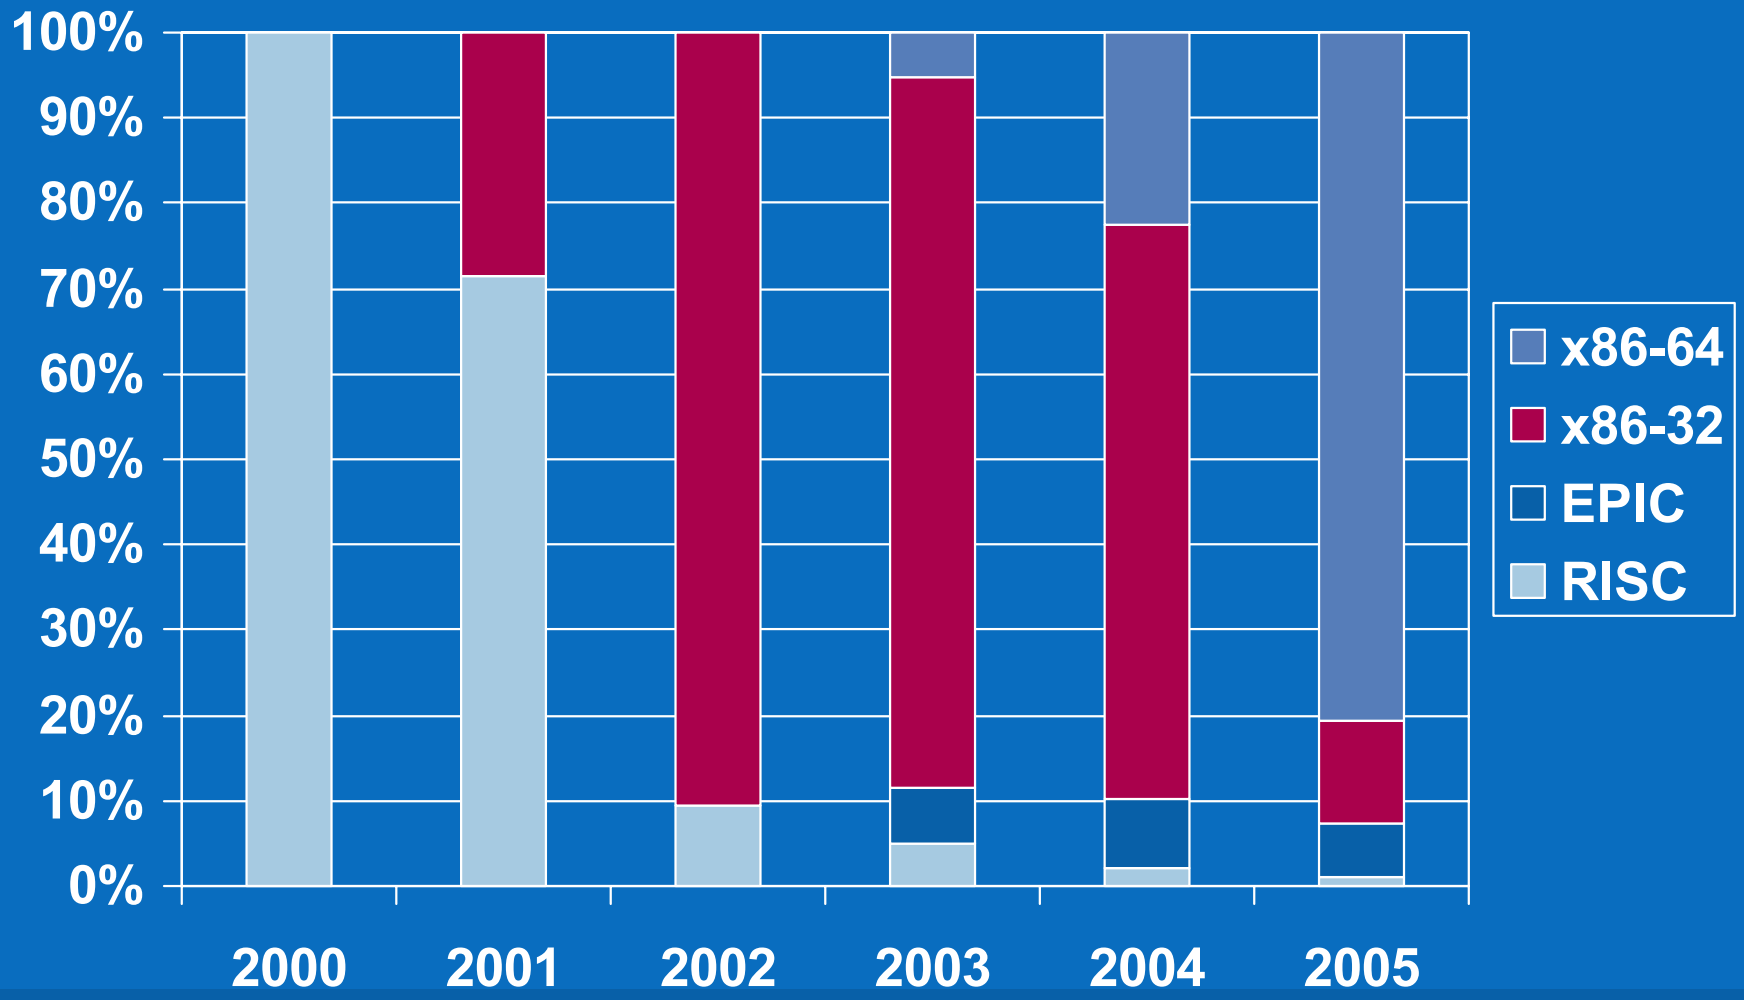

X86-64 held a 81% share in 2005

intel

Why is HPC important

- Growing market segment
- Public sector able to command (very) low margins
- Important to us as
    1. Influencer on private sector technical computing and Enterprise
    2. Early adopter of new technologies

Why Multicore

Historical Driving Forces

Increased Performance via Increased Frequency

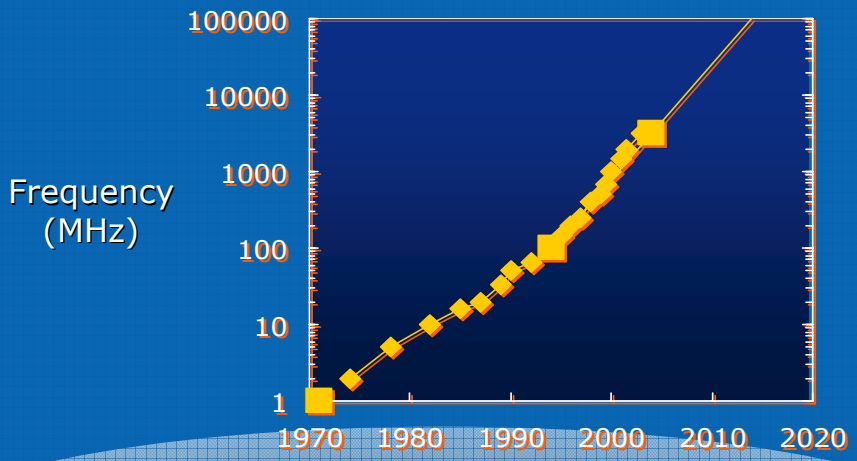

Shrinking Geometry

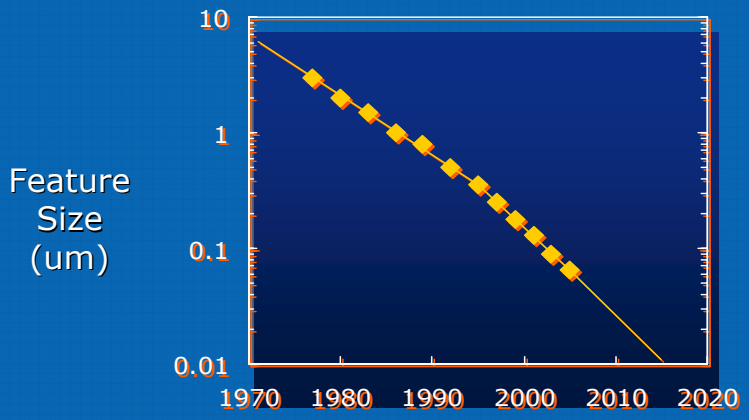

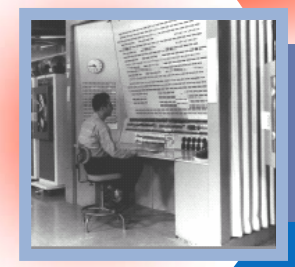

1946
20 Numbers in Main Memory

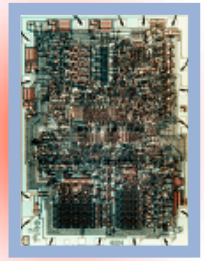

1971
I4004 Processor 2300 Transistors

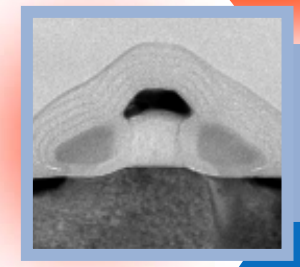

2005
65nm 1B+ Transistors

The Challenges

Power Limitations

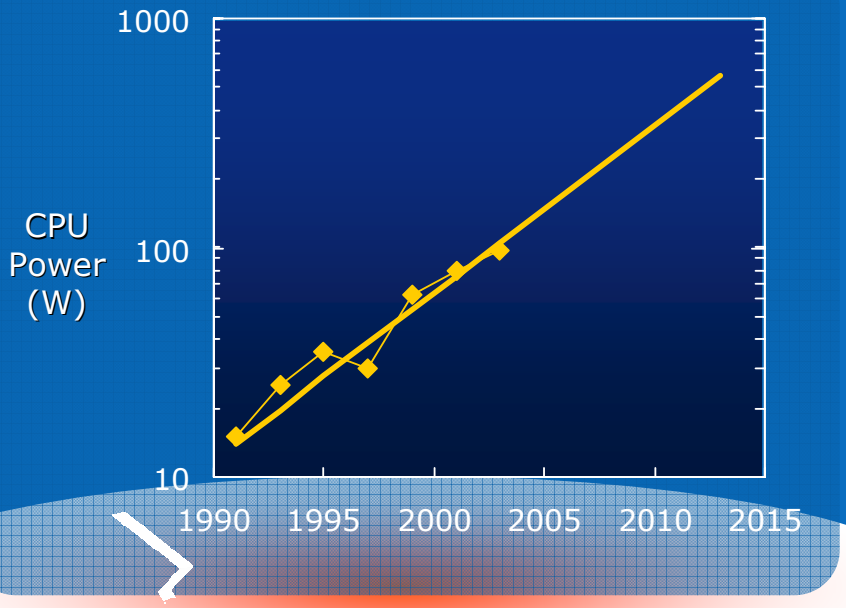

Diminishing Voltage Scaling

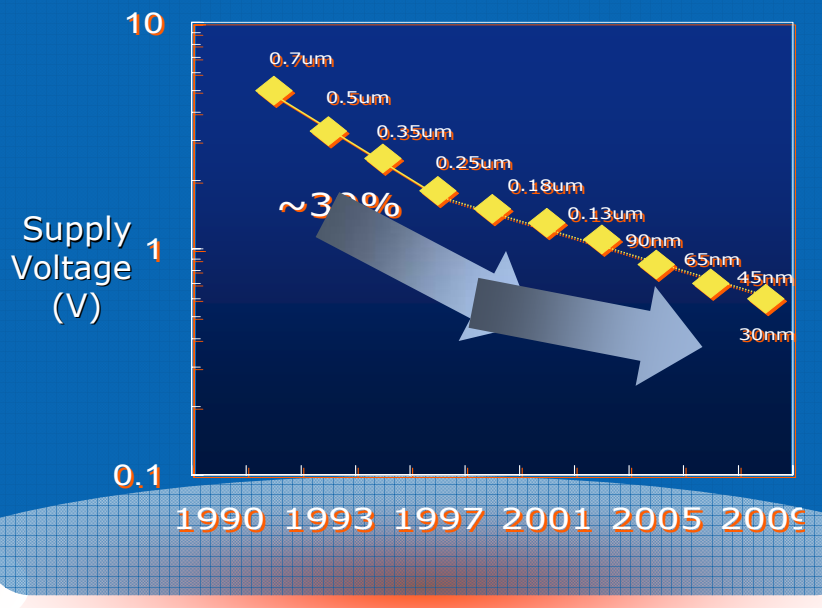

Power = Capacitance x Voltage^2 x Frequency
also
Power ~ Voltage^3

A New Era…

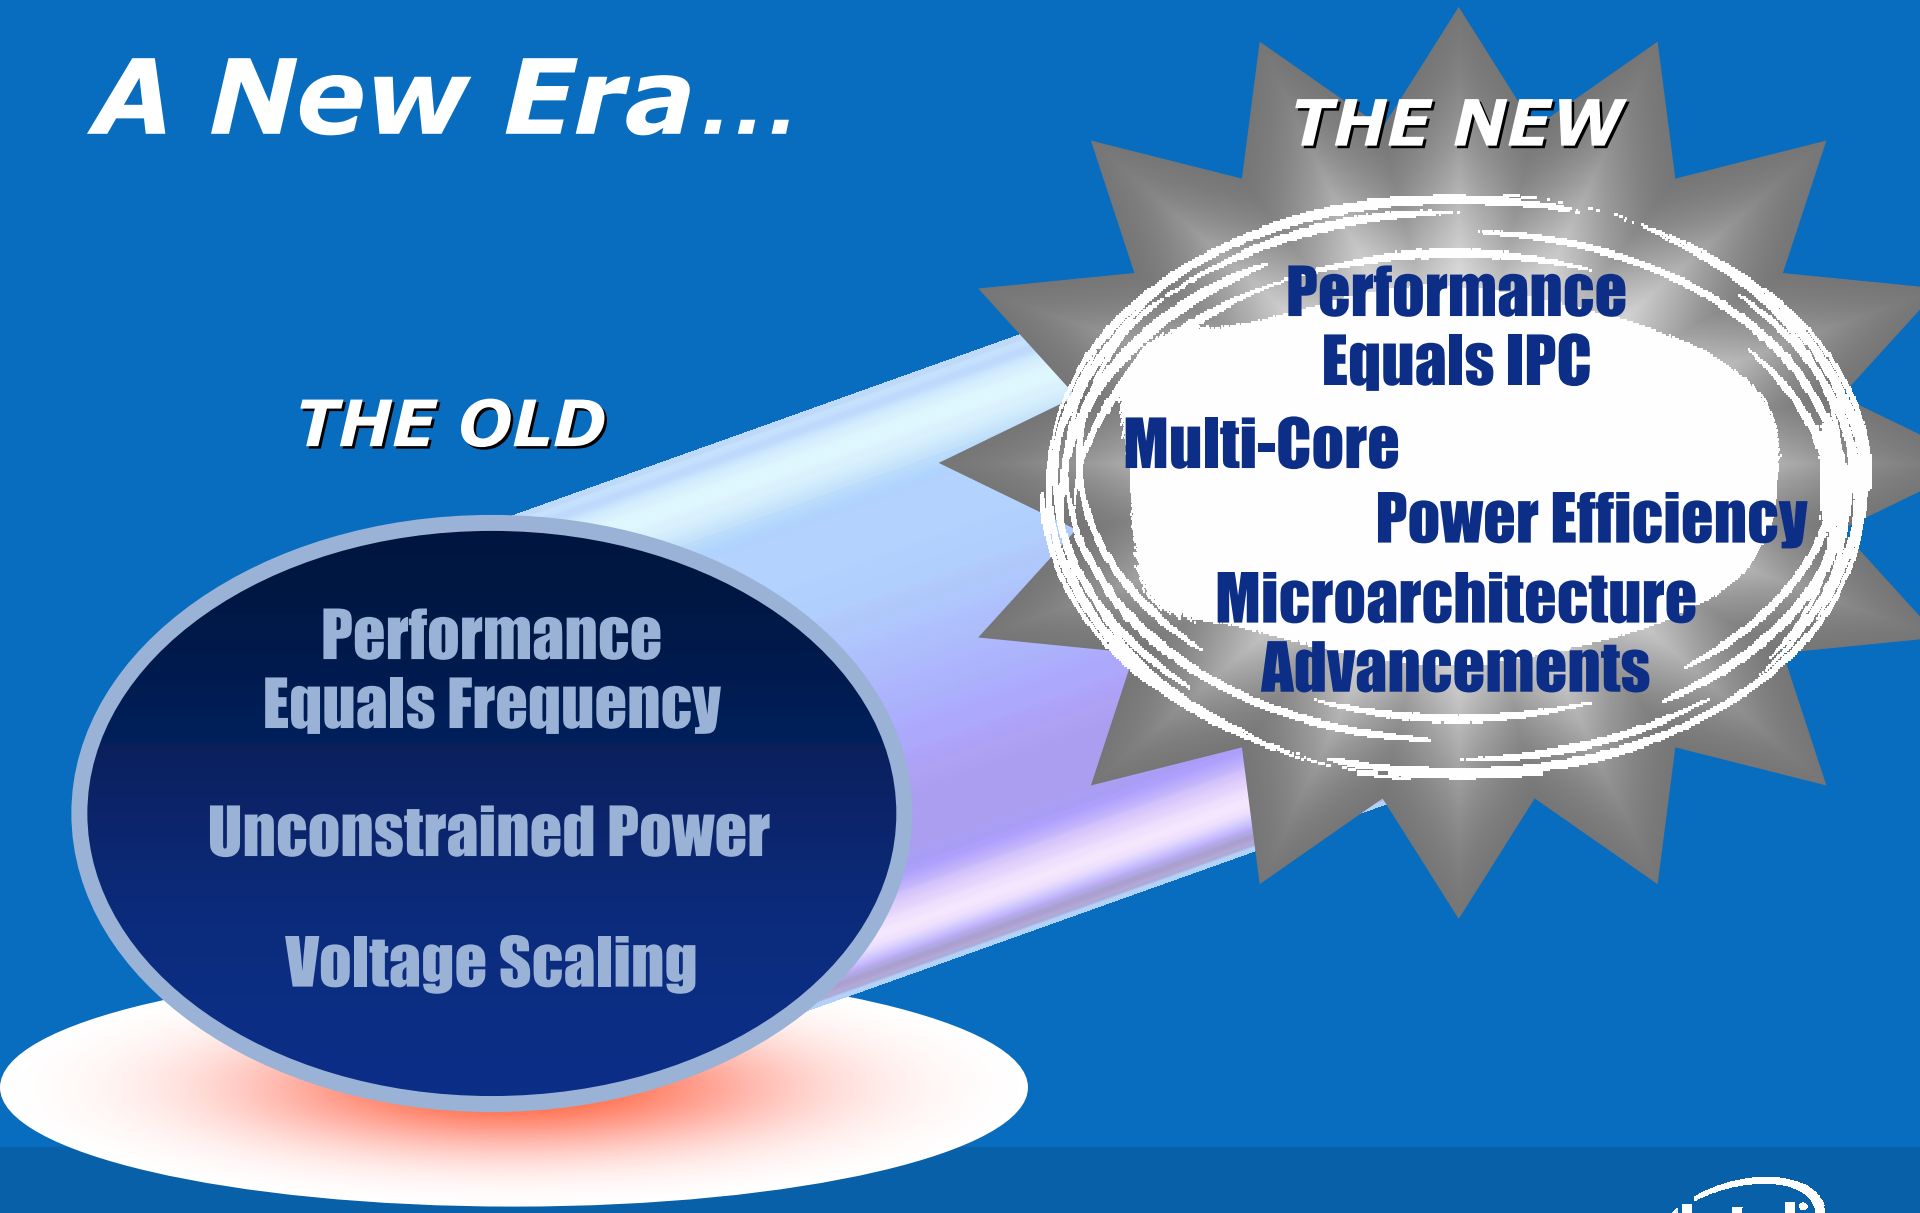
THE NEW
Performance Equals IPC
Multi-Core
Power Efficiency
Microarchitecture Advancements
THE OLD
Performance Equals Frequency
Unconstrained Power
Voltage Scaling

40 Years of Moore's Law

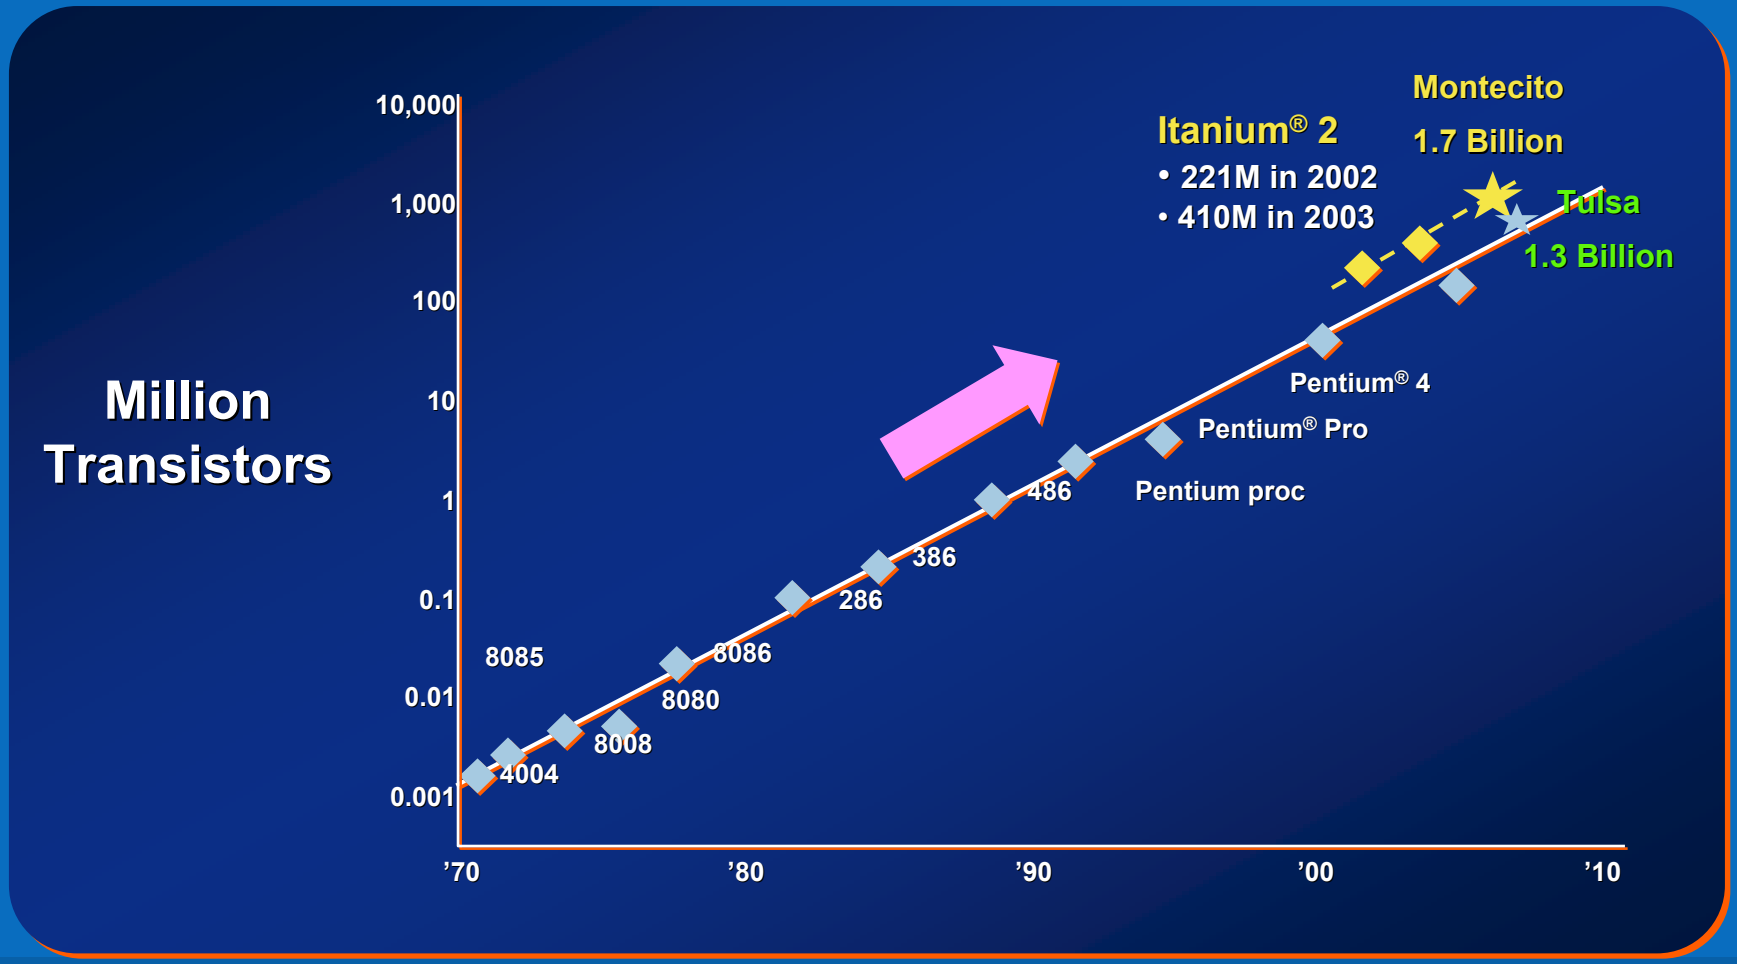

More than 1 Billion Transistors in 2005!

The Magic of Voltage Scaling

Power increases/decreases as Voltage^3

- Power = Capacitance * Voltage^2 * Frequency
- Frequency α Voltage in region of interest

Good news is that voltage scaling works

10% reduction in voltage yields:

- ~10% reduction in frequency
- ~30% reduction in power
- Typically ~7-8% reduction in performance

Voltage scale to maximize perf in fixed power

*may not be feasible

Simple Dual Core Example

Assume Single Core processor at 100W

- 80W for core, 20W for cache and I/O

Dual core within ***same*** power envelope

- 20W for I/O and cache
- 40W per core

- Reduce voltage by ~20% to reduce core power to 40W
- Frequency reduces by ~20%
- Single thread perf reduces by ~15%
  - Several techniques to help reduce this loss in the future
- Throughput increases by 70-80%

Why Multi-Core?

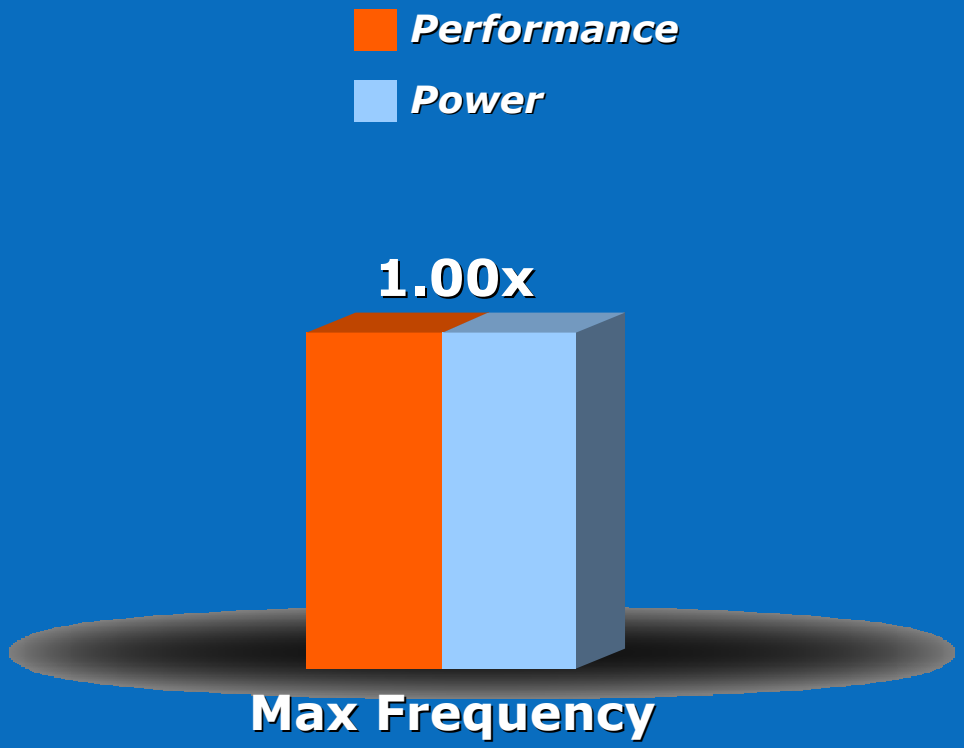

Relative single-core frequency & Vcc

Source: Intel Jun 2006

Over-Clocking Example

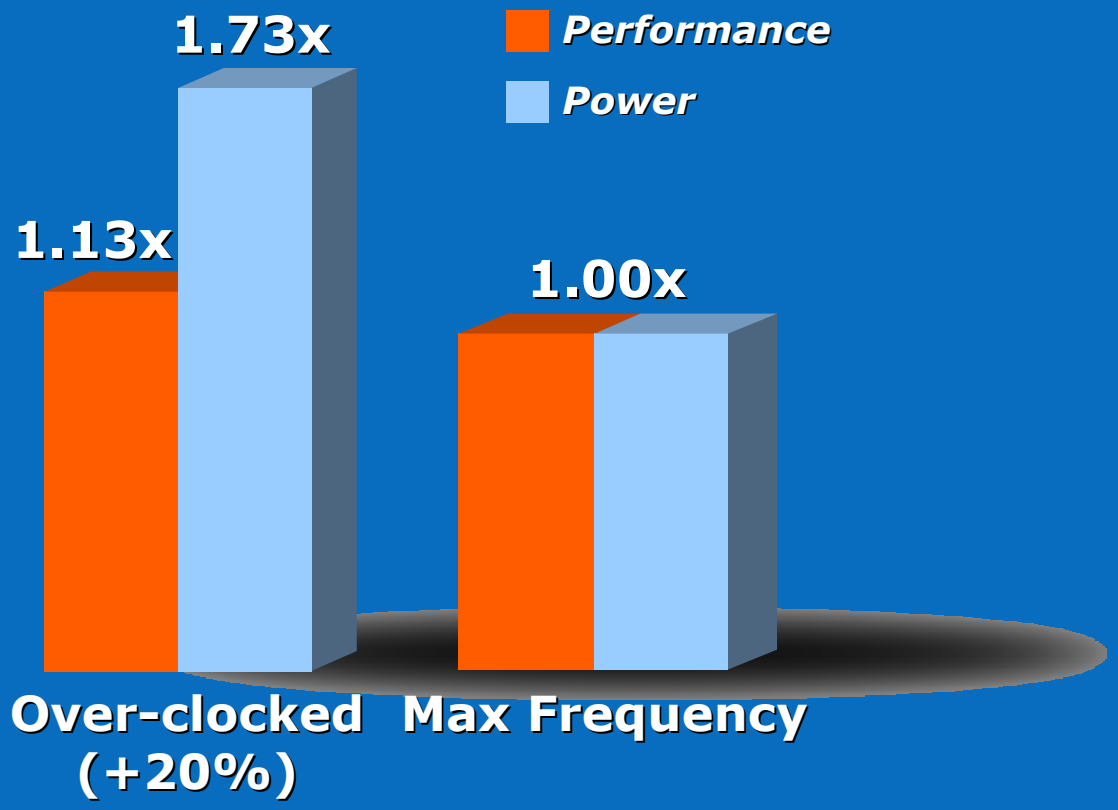

Relative single-core frequency & Vcc

Source: Intel Jun 2006

Under-Clocking Example

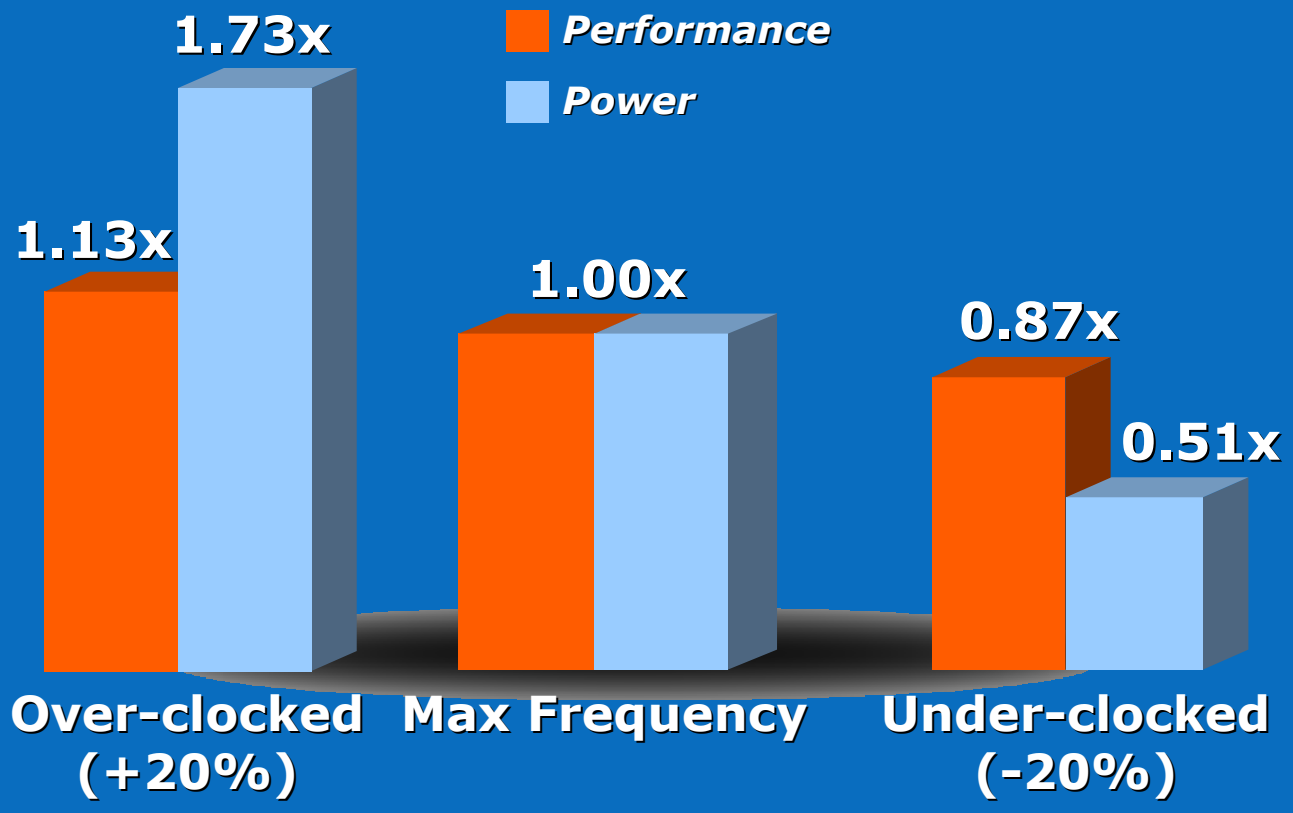

Relative single-core frequency & Vcc

Source: Intel 2006

Multi-core: Energy-Efficient Performance

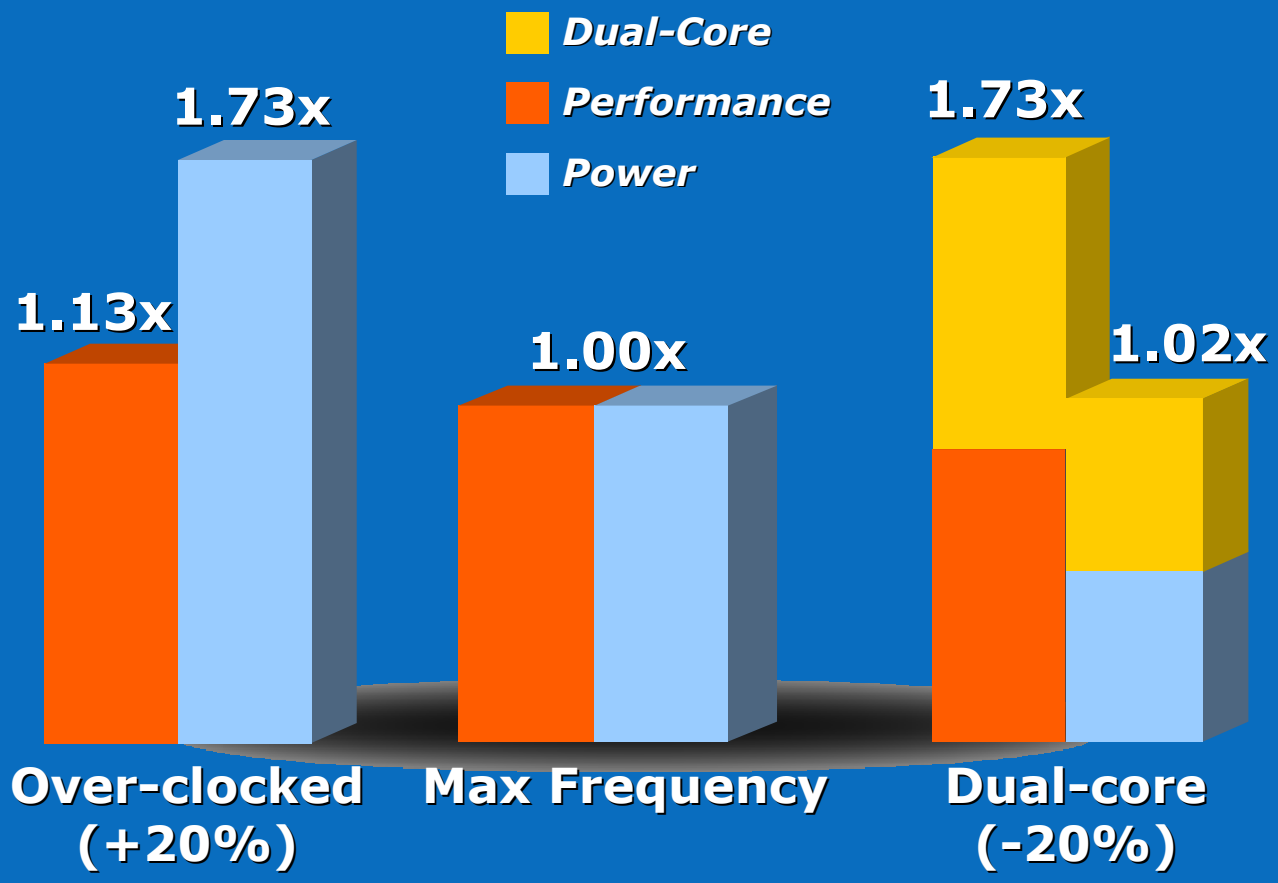

Relative single-core frequency & Vcc

Source: Intel 2006

Intel® Core™ Microarchitecture

The Energy Efficient Performance Leader

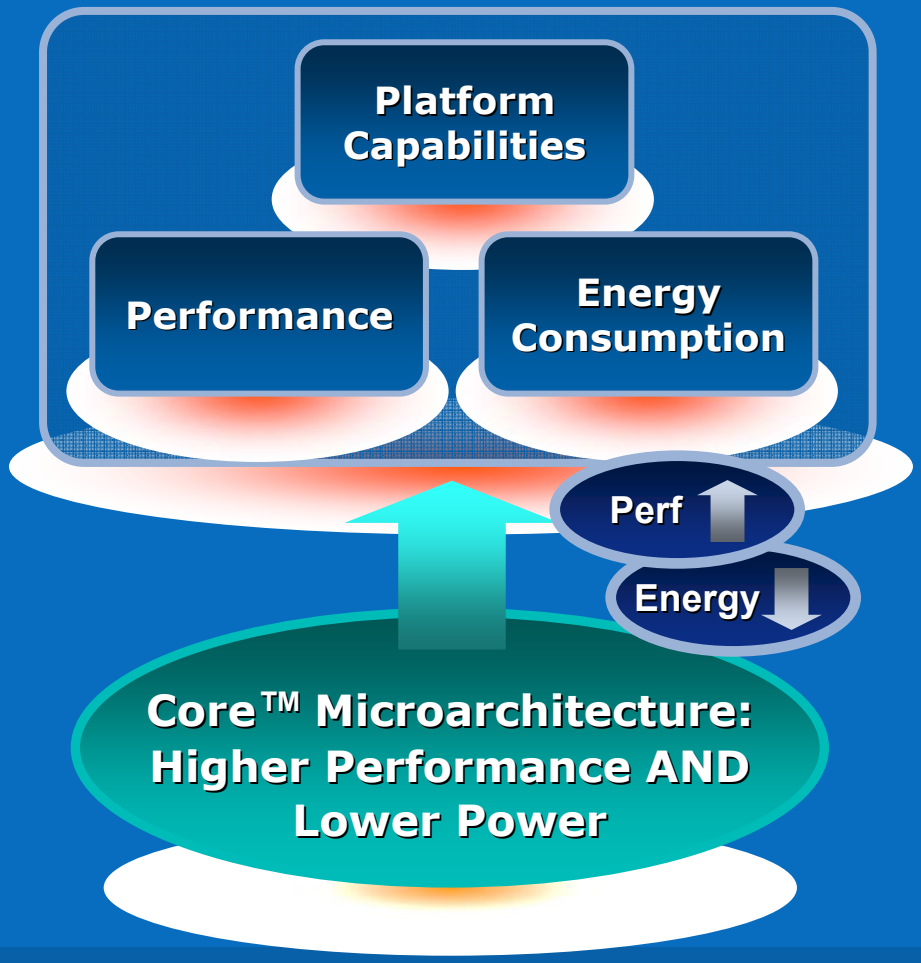

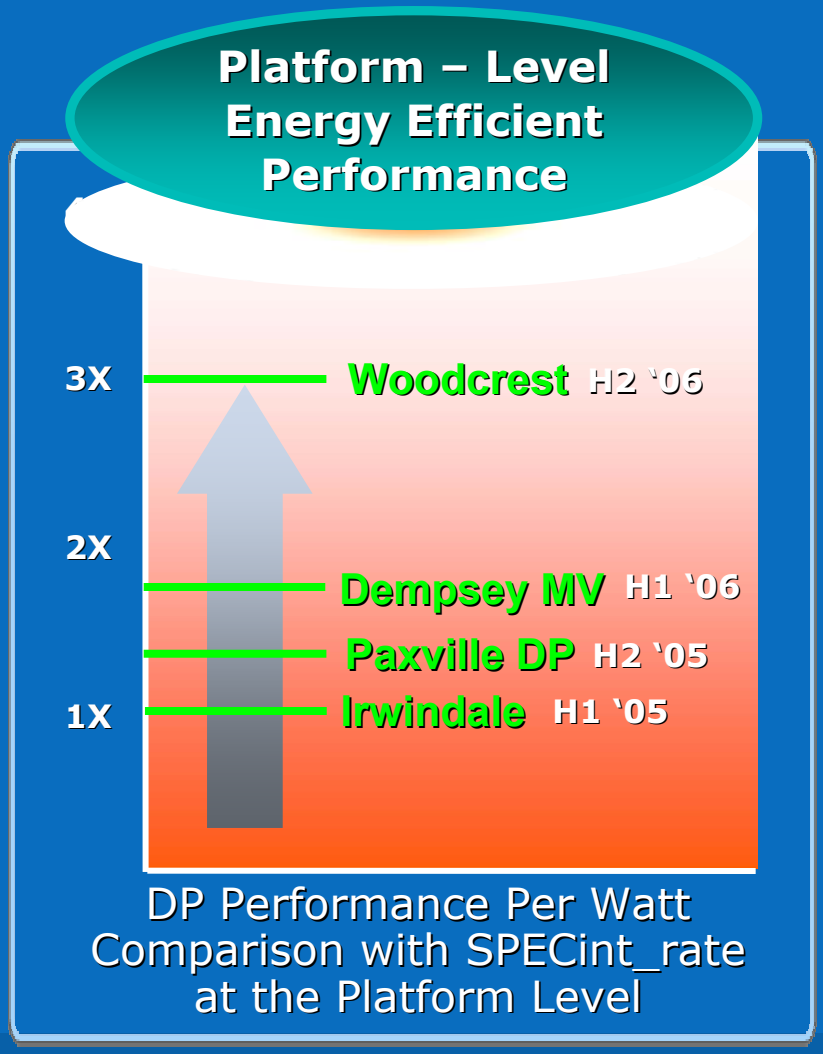

Source: Intel®

Intel® Core™ Microarchitecture

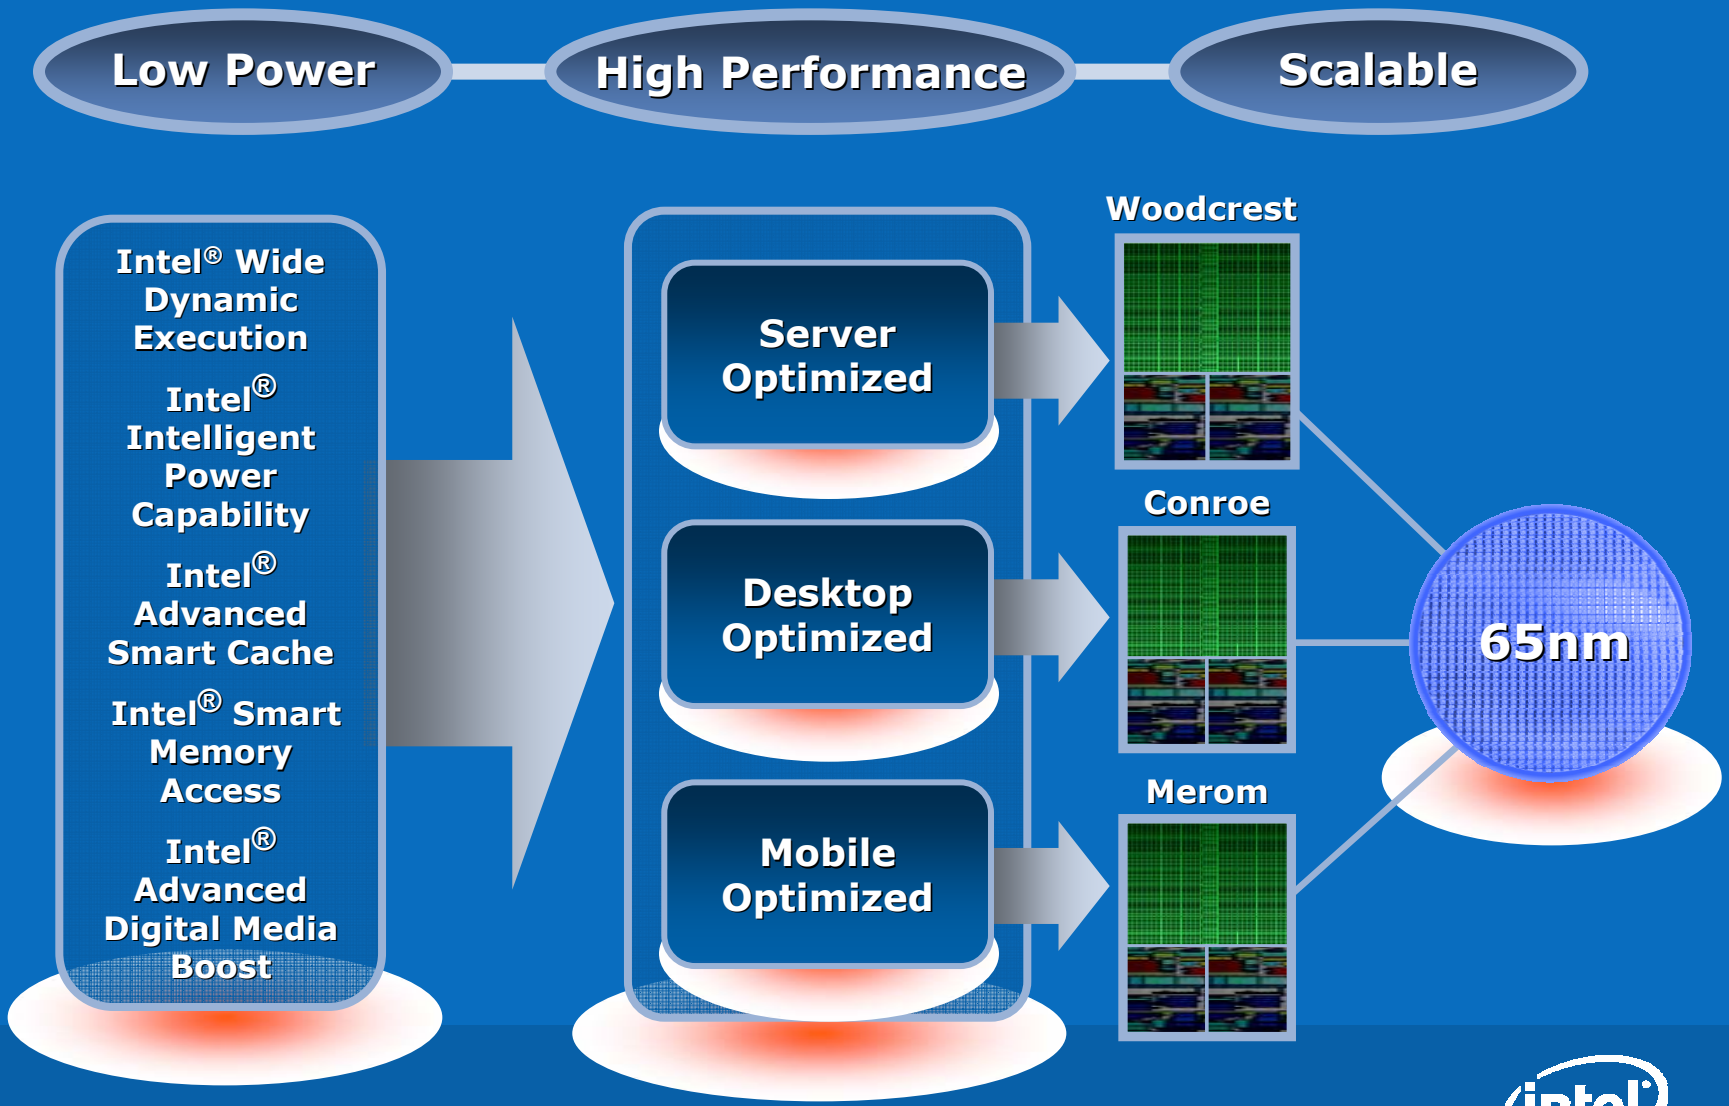

*Graphics not representative of actual die photo or relative size

HPC at Intel

HPC Strategy

Top 10% of Top 500

- High touch
- System design applies to broader server technology roadmaps

Next 90% of Top 500

- Scale through OEM support programs
- Reactive marketing support
- Promote key wins and PetaScale efforts as HPC halo for Top 500 influencing

Volume HPC Markets (90% of TAM)

- ISV/IHV co-marketing plans and programs for priority commercial verticals
  - Manufacturing, Life Sciences, Energy, Finance
- Connect server platform technology demand creation programs to Grid/Cluster solutions in volume HPC/HDC
- Linux cluster architecture recipe development and ISV validation program
- Microsoft CCS co-marketing program

Architecture focus by vertical

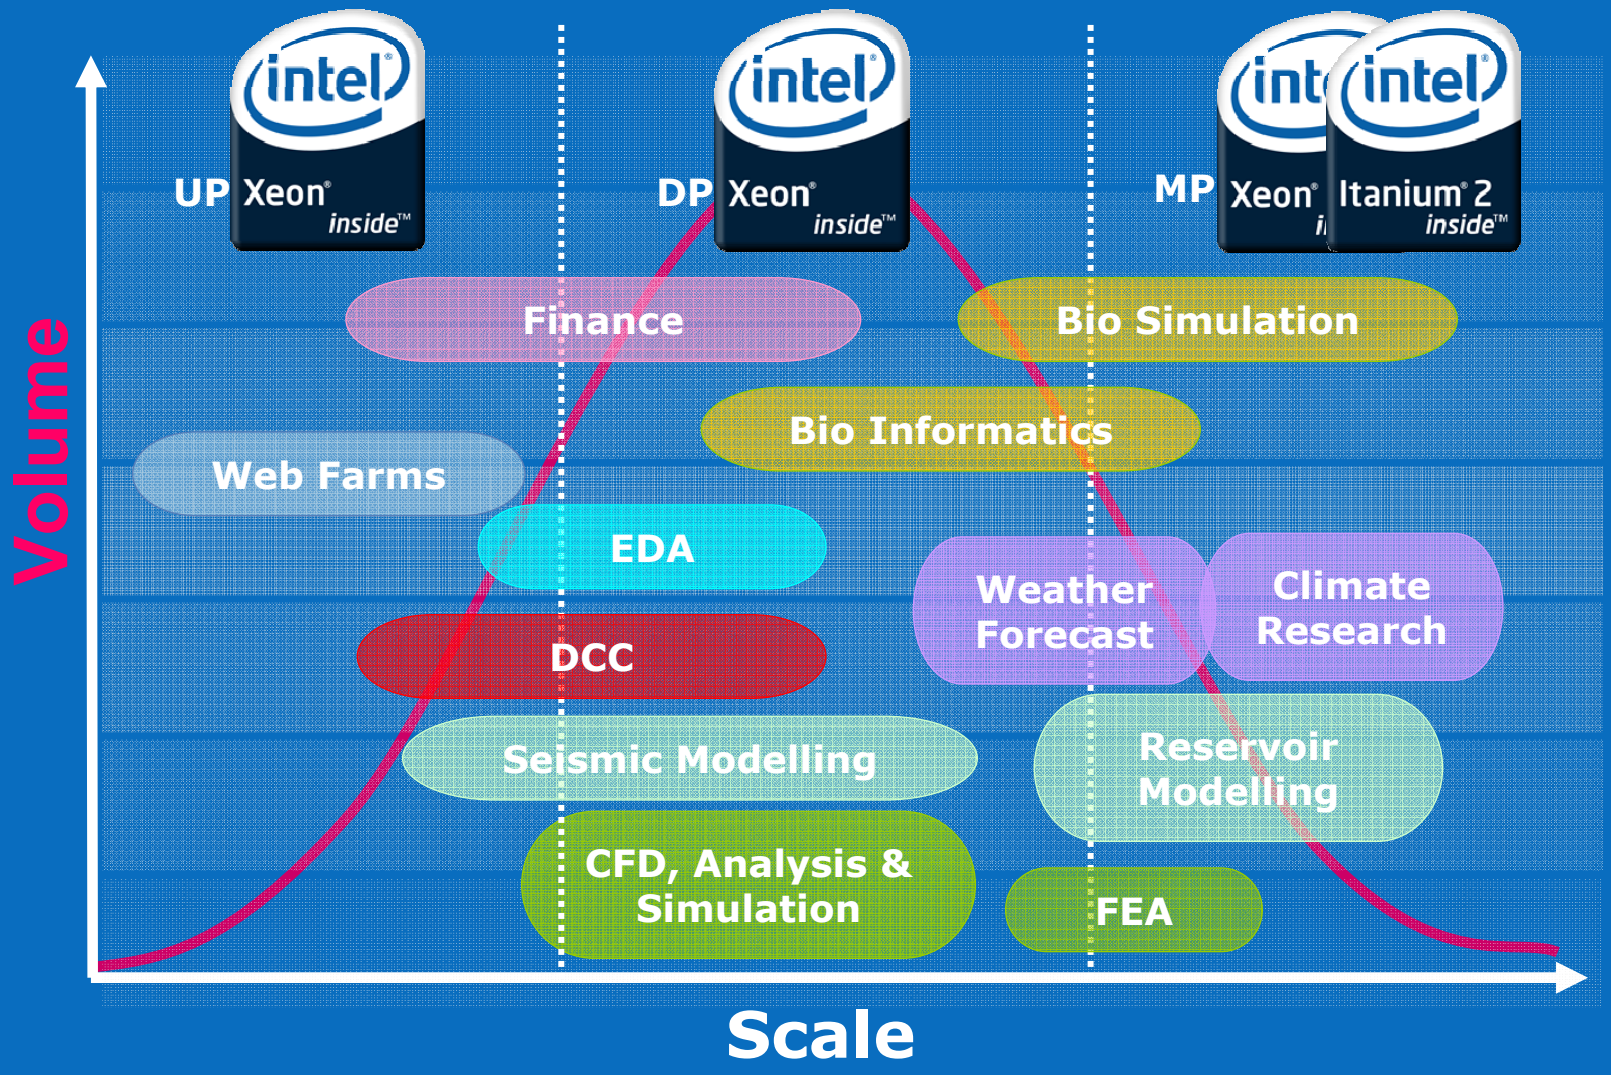

* vertical's placement not indicative of volume

HPC Product Segmentation

High density, data intensive thin node (Bus Guzzler Apps)

- High density, targeted feature set
- Balances processor, memory and network BW
- Cost sensitive

UP

High density, compute intensive fat node (Bus Miser Apps)

- Broader feature set
- Greater memory capacity
- RAS capability
- Less price sensitive

DP

Guzzlers accounts for more than half of the market

Target UP at Personal, Workgroup clusters, low end, highest revenue growth

HPC/HDC Cluster Segmentation Model

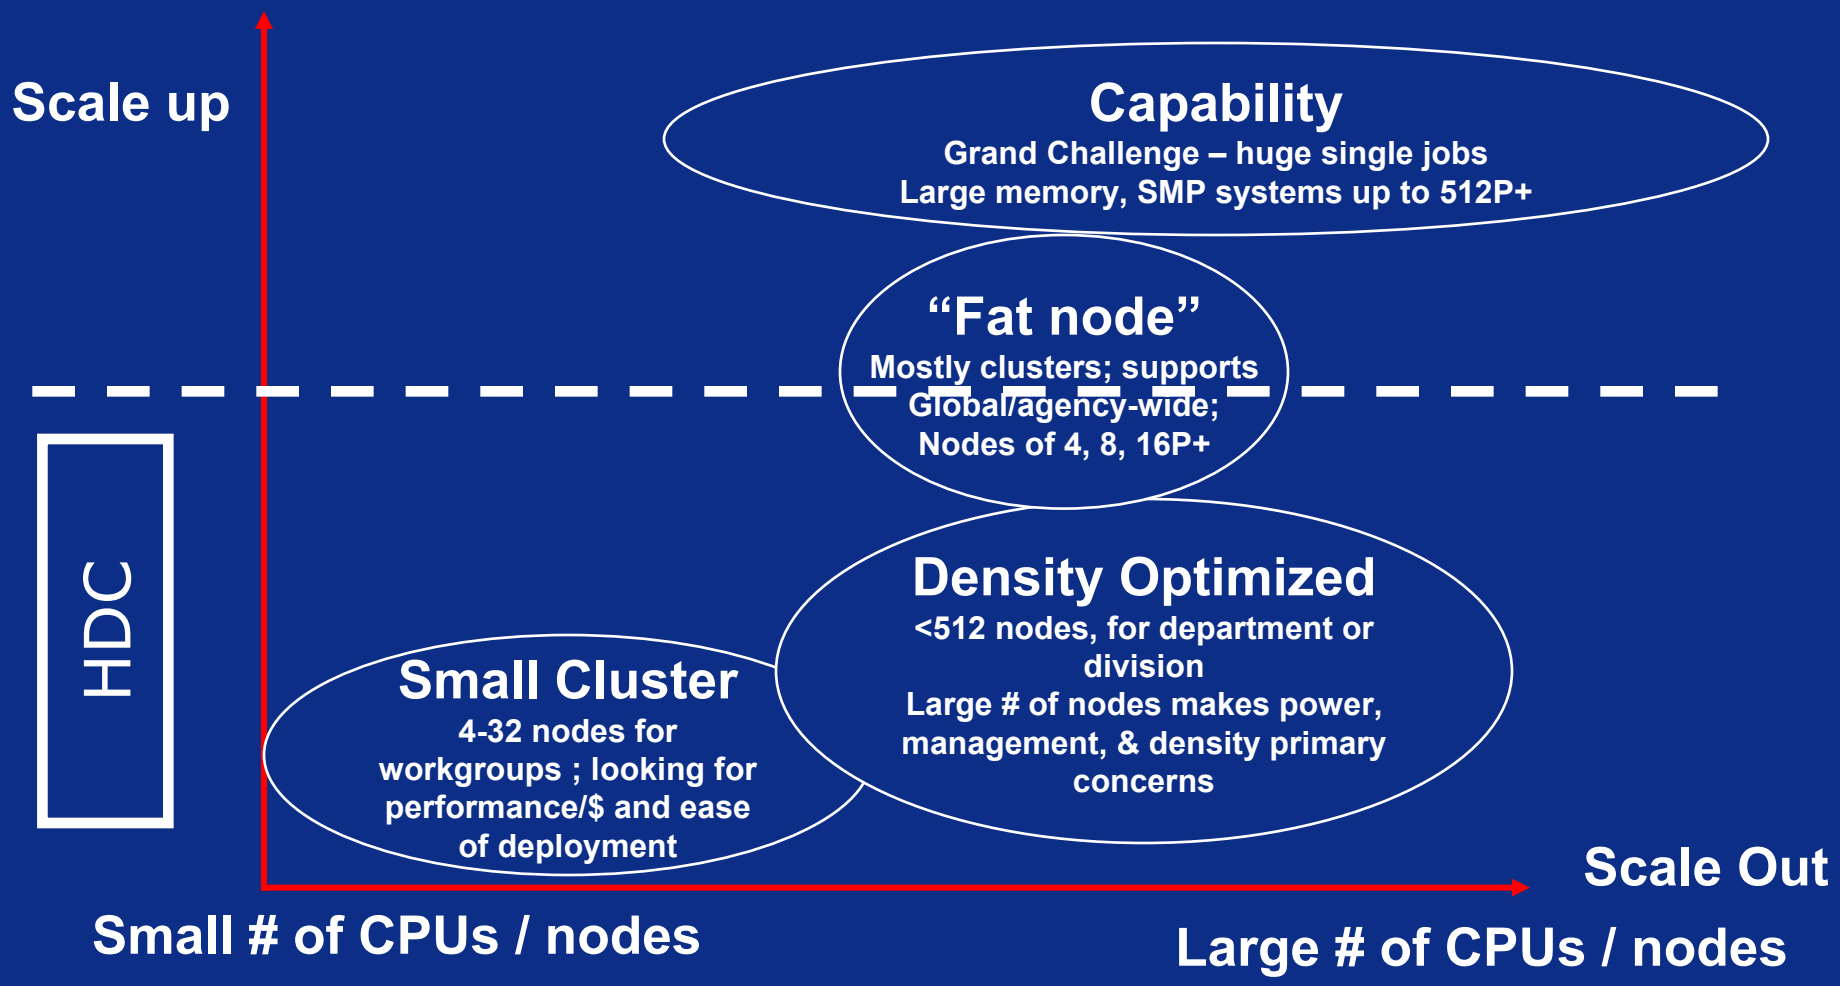

HDC - the largest TAM and Revenue Segments, and highest CAGRs

Towards a

Petascale Machine

Processor Performance

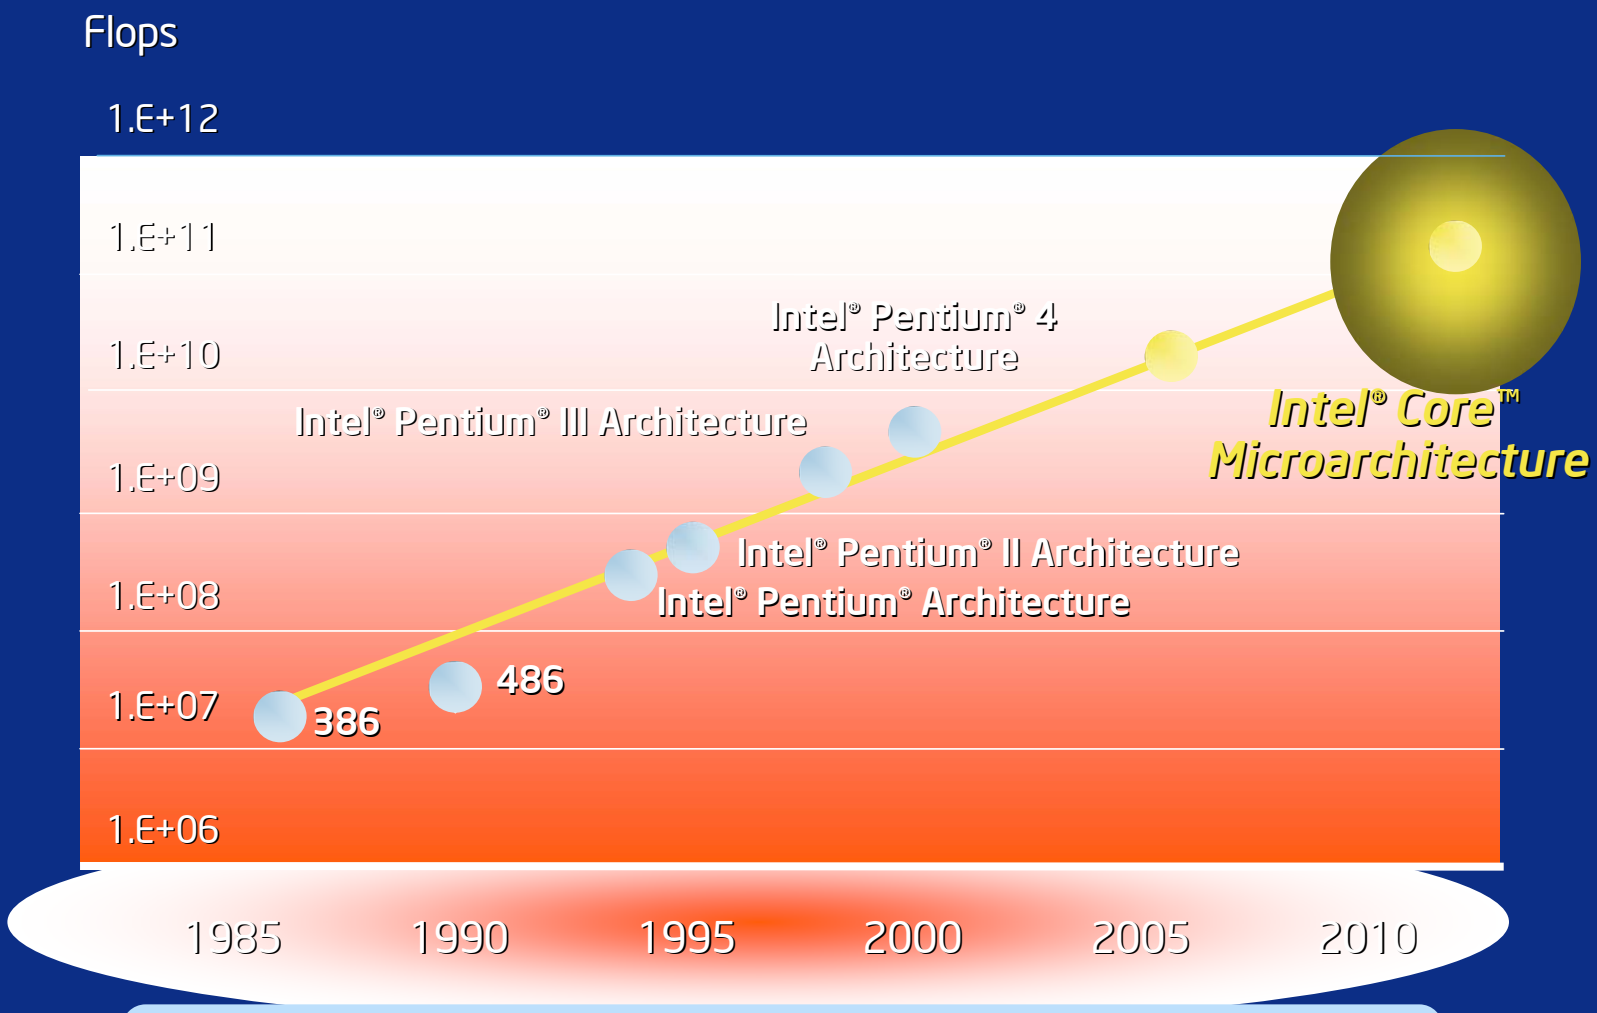

Reaching Petascale with ~100,000 Processors in 2010*

Assuming approx. 100Gflops processors

* Petascale assumes 10's of PF Peak Performance and 1PF Sustained Performance on HPC Applications.

Multi-threaded Cores

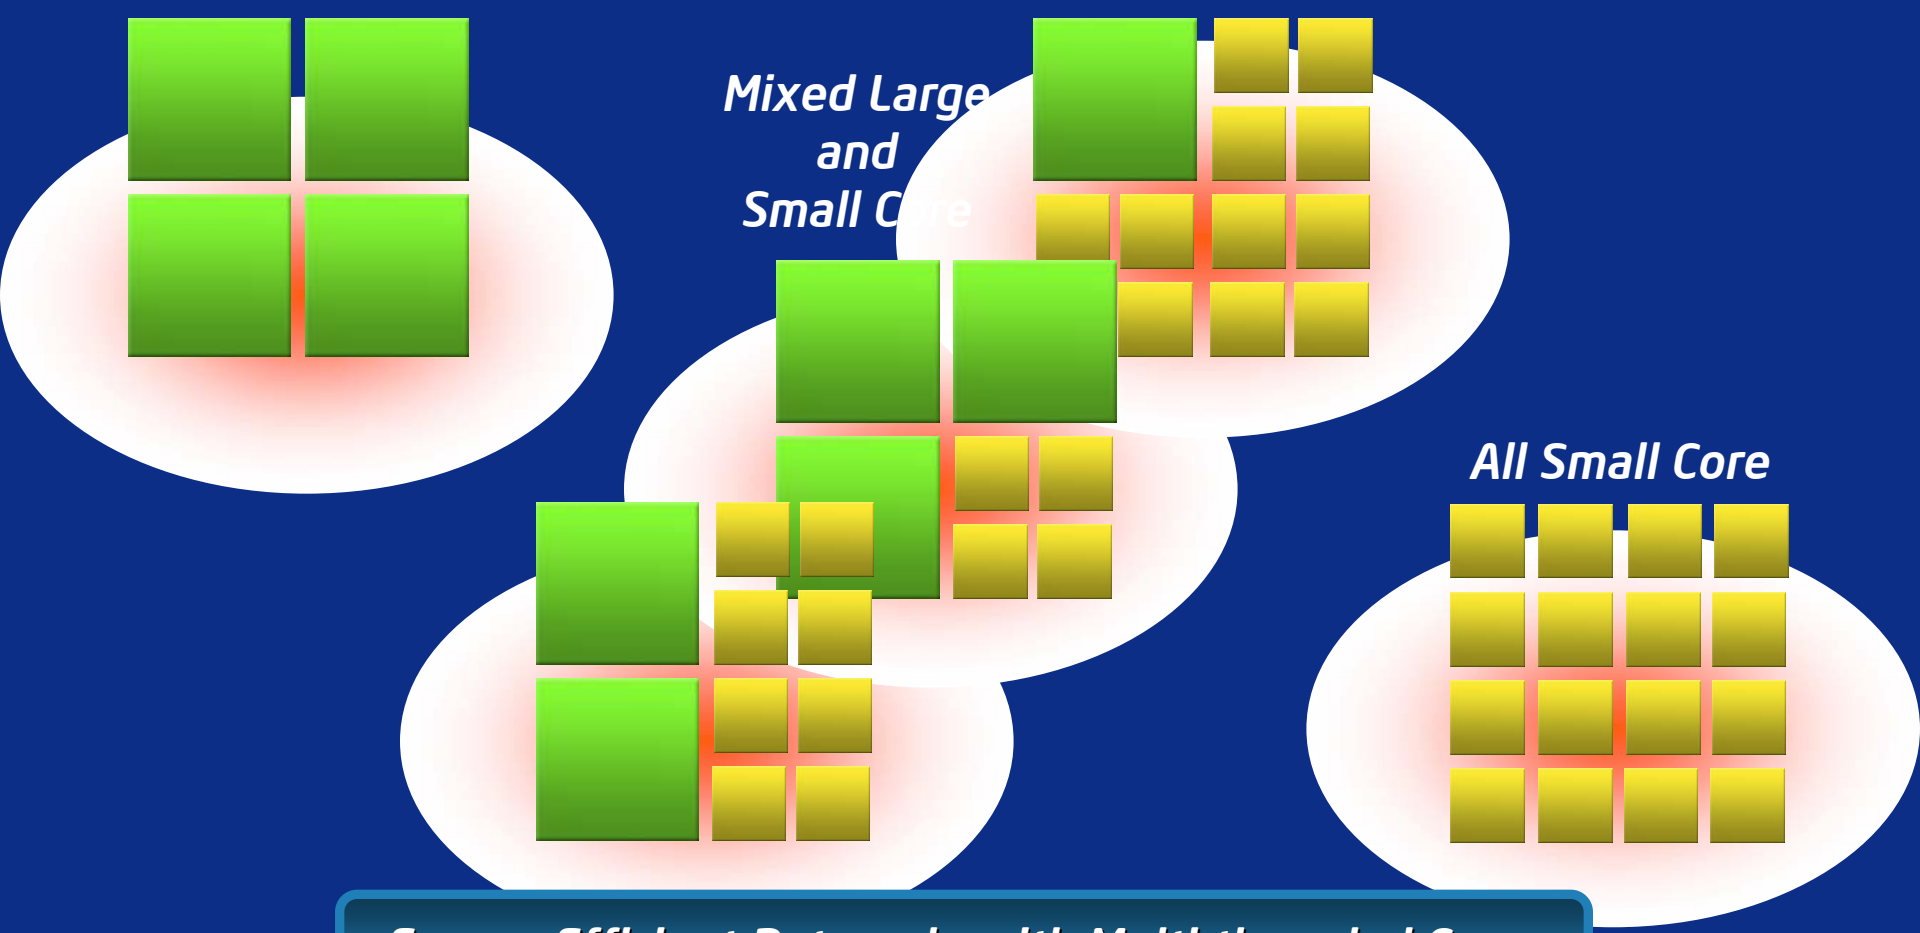

Energy Efficient Petascale with Multi-threaded Cores

Note: the above pictures don't represent any current or future Intel products

intel

Increasing Throughput through Parallelism

Amdahl's Law: Parallel Speedup = 1/(Serial% + (1-Serial%)/N*)

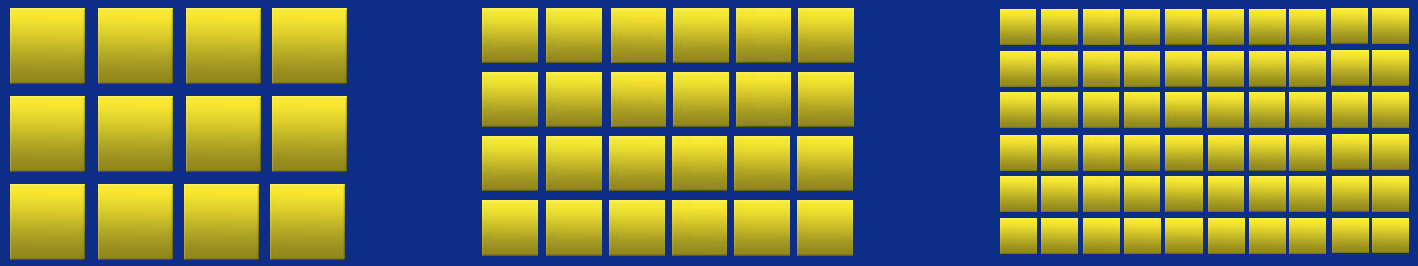

Single Core Performance

Relative Performance

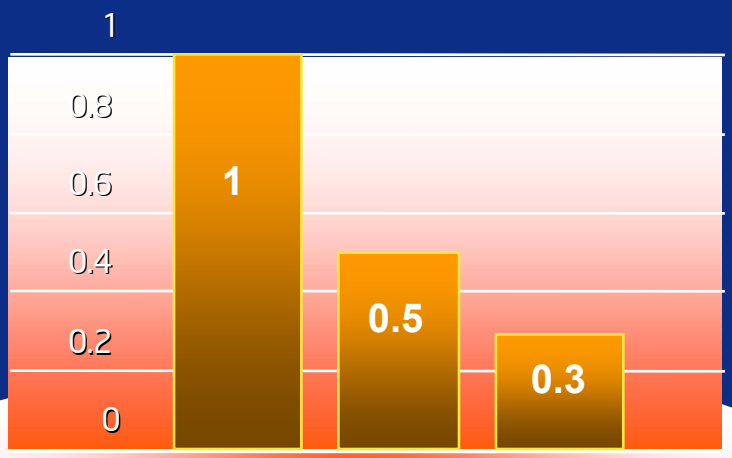

System Performance

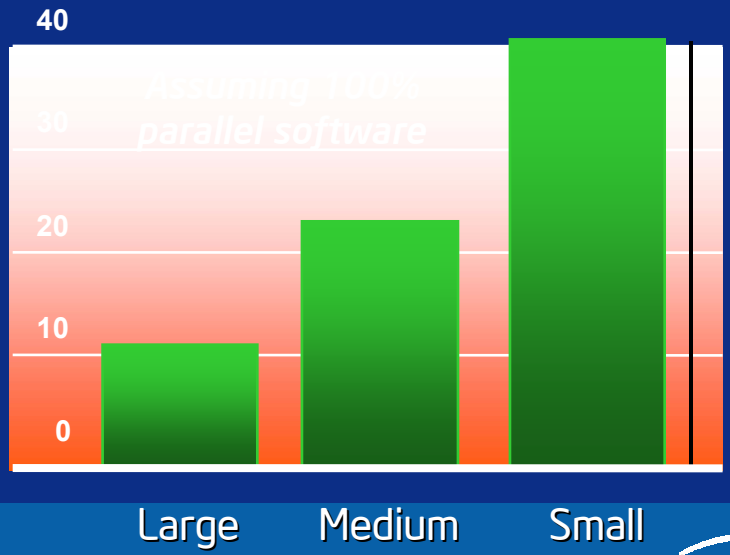

* N = number of cores

Increasing Processor Performance

Through Multi-threaded Cores

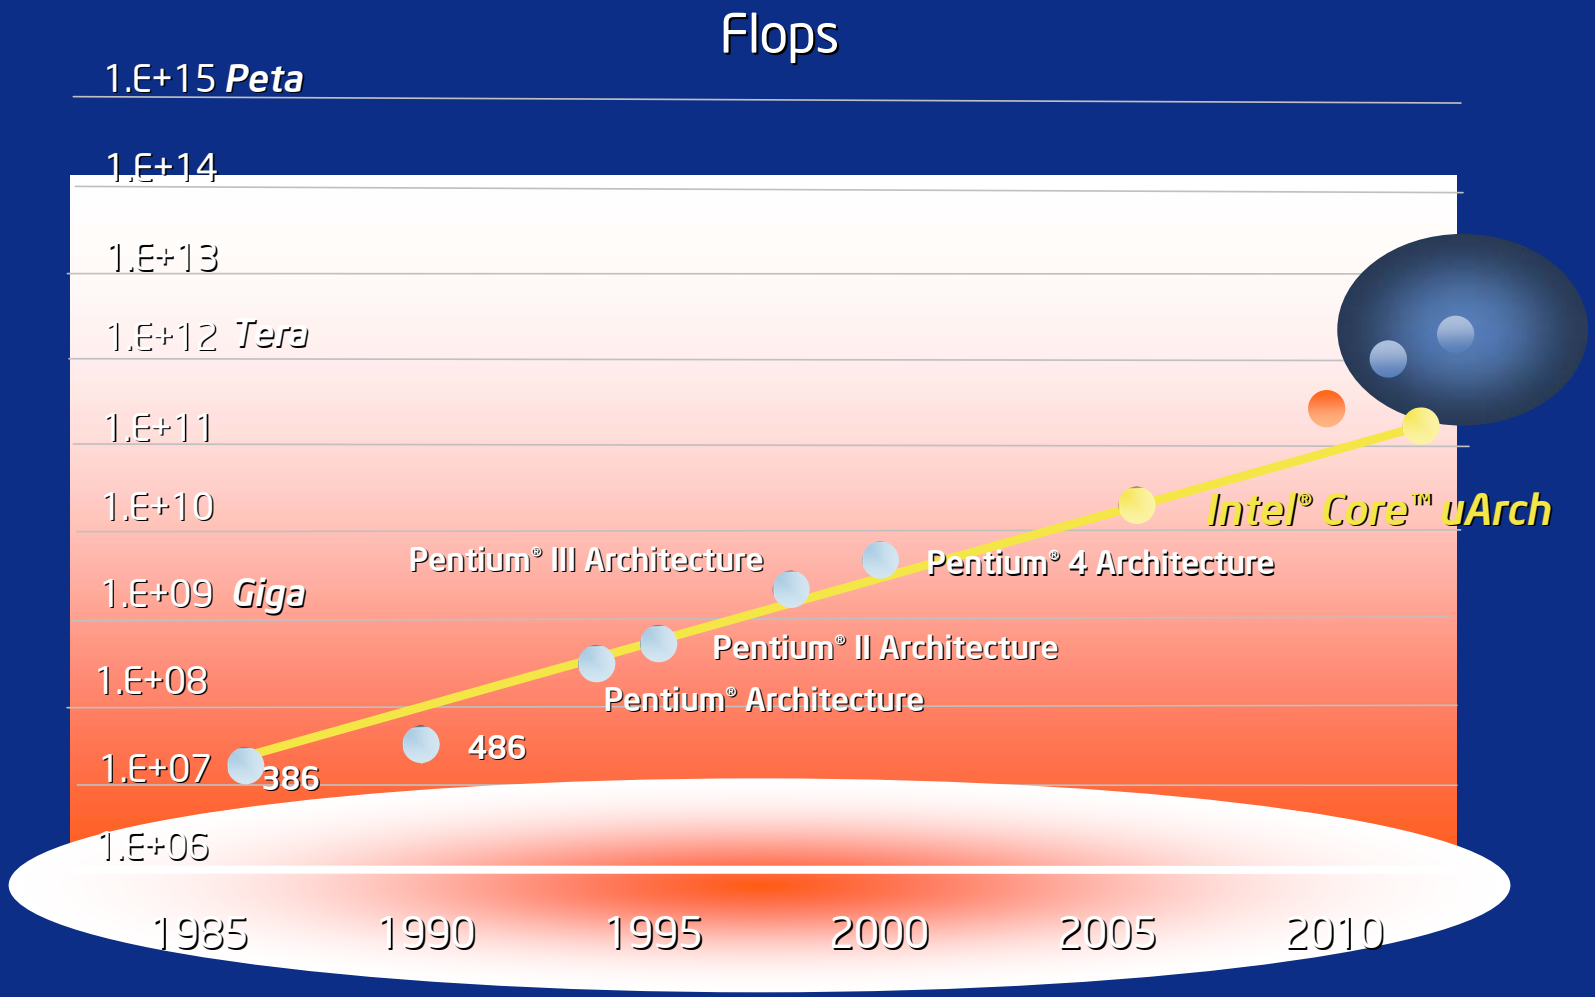

Reaching Petascale with ~5,000 Processors

Increasing Memory Bandwidth
to Keep Pace

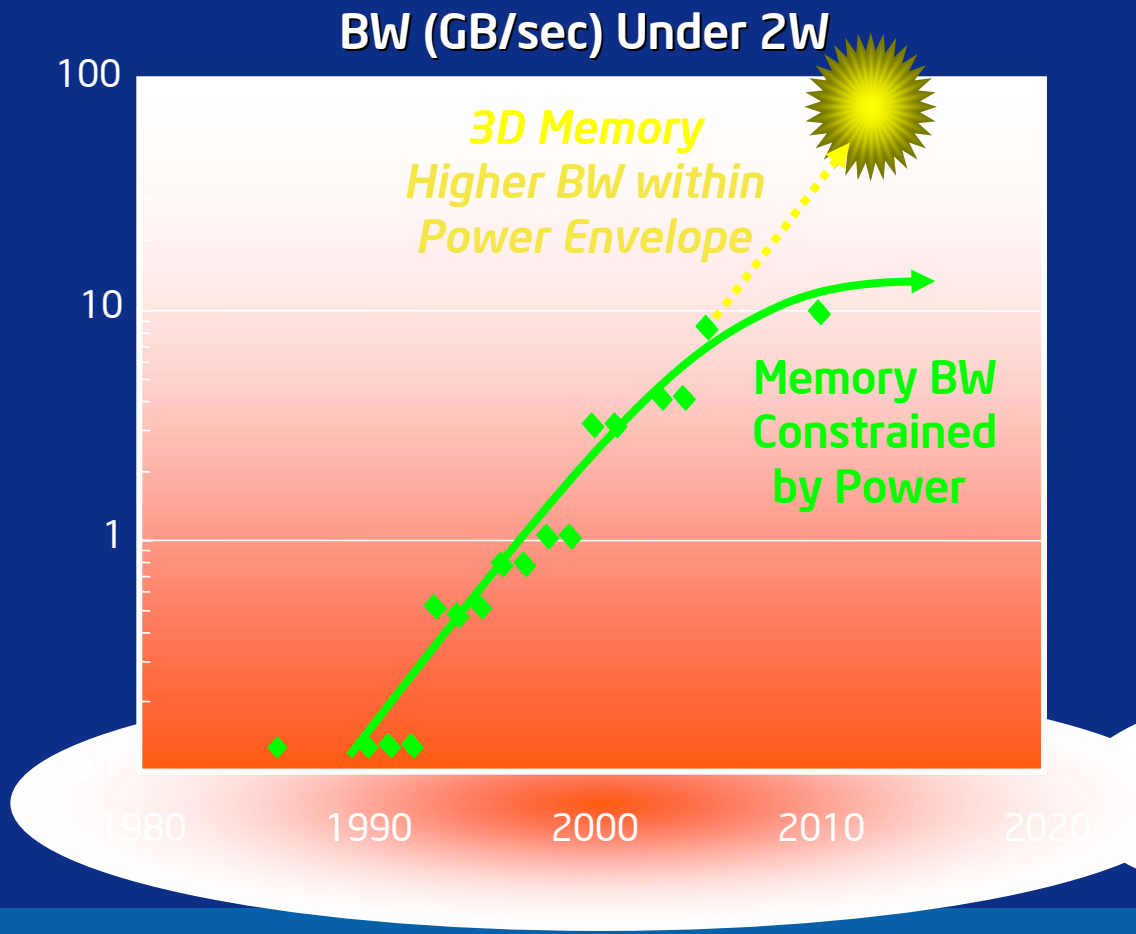

3D Memory Stacking

Power and IO Signals Go Through DRAM to CPU

Thin DRAM Die

Through DRAM Vias

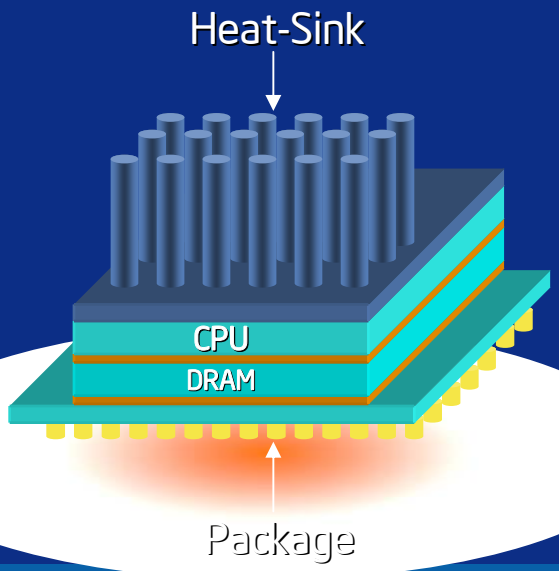

Source: Intel

Implications

The Top500: Reaching Petascale

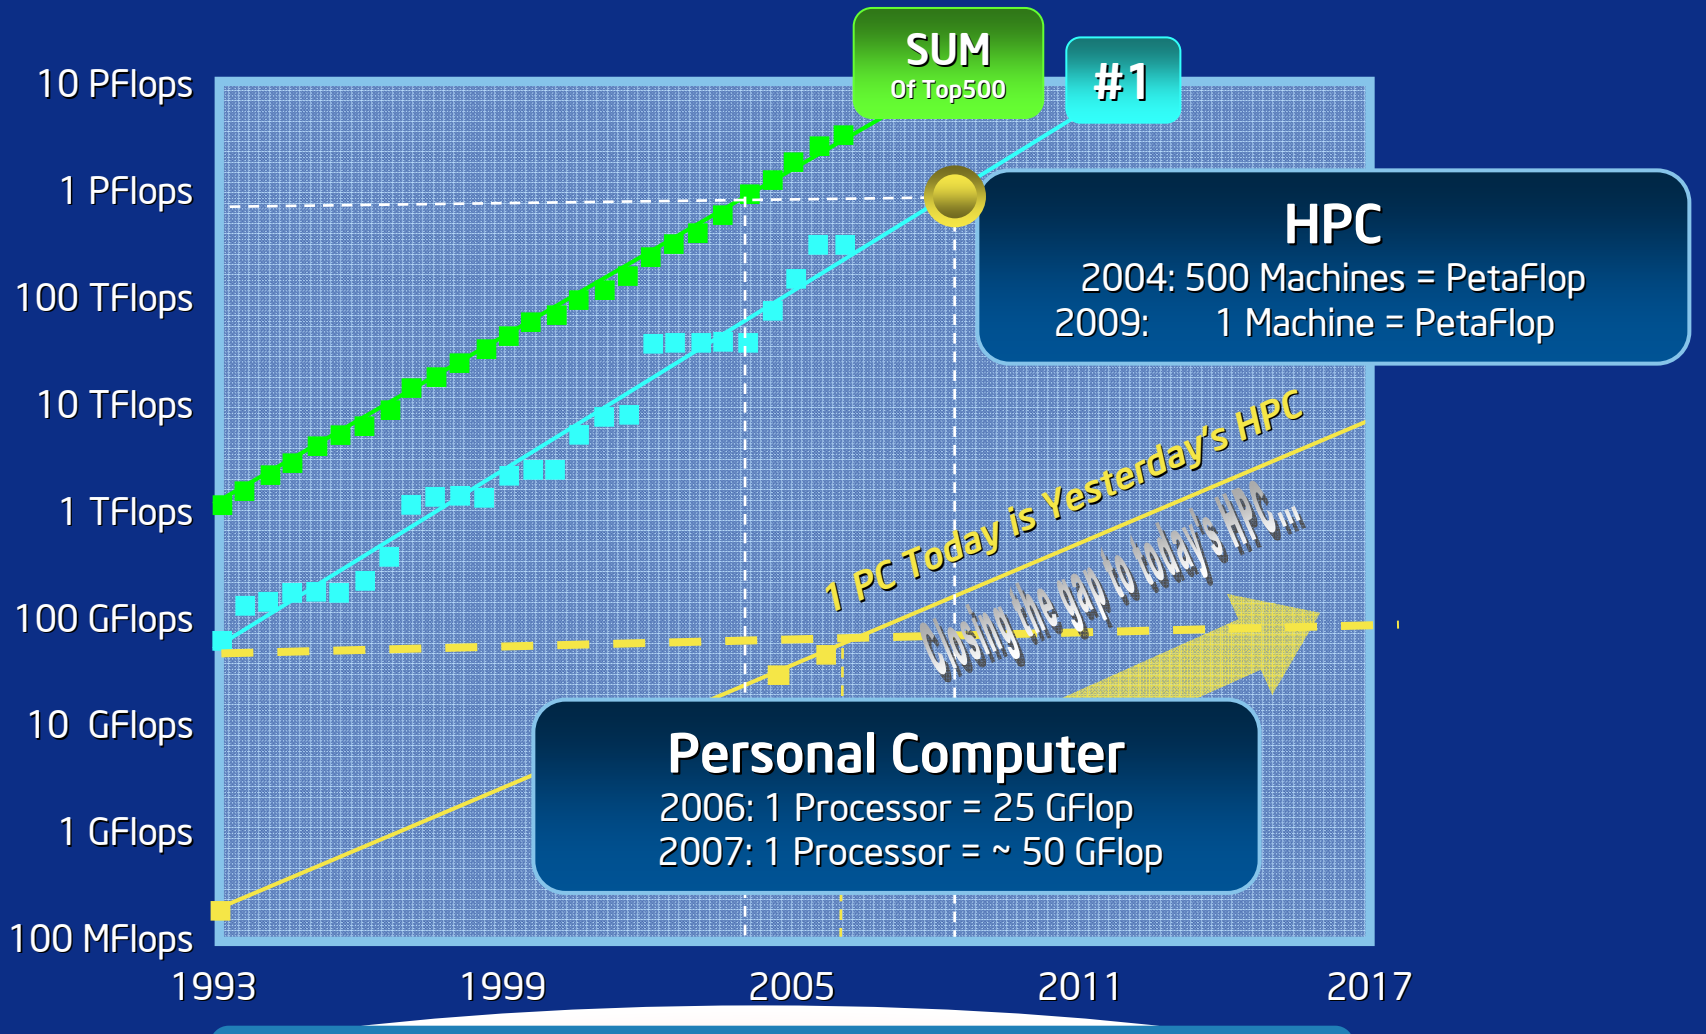

2008: Peak and Linpack PetaFlop – 2011: Sustained

Source: HPC - www.top500.org, June 2006; PC - Intel

Pace of advance of supercomputing

- Megaflop range – CDC 7600 – early 1970's
- 1GF – 1980 – Cray-XMP
- 1TF - 1996 – ASCI Red, Intel
- 1PF – 2008 .. By ??

1,000,000x increased peak performance in <30 years!

What about parallelism?

- The 1980 1GF system:
  - 4 vector processors, 64 elements vectors → 256-way para
- The 1996 1TF system:
  - ~ 9,000 processors → 9K-way para
- 2008 1PF (peak)
  - ~ 18,000, 4 core, with 4 flops/clock → ~ 288K-way para
- ~ 2010 sustained PF (10-20PF peak)
  - At least 32 concurrent ops in a socket; > 1,000,000 threads

Level of parallelism increased by O(1,000) [1980-2008]

The real big challenge

- Programming model for hierarchical expression of parallelism
  - Inside the core → socket → node → cluster [→Grid?]
- Algorithms that scale
  - Eliminate serial code
    - Even 1% serial means max of <100x speed up; what to do with > 100,000 sockets??
- Is it hopeless? – no.
  - Examples today of ~65K way scaling (on BlueGene)
  - New techniques – such as ensemble methods in weather 7 climate..

In conclusion …

Awesome opportunities

Intimidating challenges


intel®
Leap ahead™